



Administration Guide

Active Circle Storage System

Version 4.5

Active Circle

Administration Guide: Active Circle Storage System

Version 4.5

Published September 2016

Revision 3.0

Copyright © 2018 Active Circle. All rights reserved.

The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means — electronic, mechanical, magnetic, optical, chemical, manual, or otherwise — without the prior written permission of Active Circle. The Information in this document is subject to change without notice.

Active Circle is a registered trademark of Active Circle SA.

All other marks and names mentioned herein may be trademarks of their respective companies.

Active Circle S.A.

26, rue du Faubourg Poissonnière

75010 Paris

FRANCE

Tel: +33 1 34 65 90 29

www.active-circle.com

Contents

Introduction	xi
Part I. System Overview	
1. What's New	1
2. Features and Concepts	3
Main Features	3
Key Concepts	4
3. The Active Circle Service	9
Starting the Active Circle Service	9
Stopping the Active Circle Service	9
Active Circle Service Extension	9
Specifications	10
Implementation	11
4. The Administration Tool	13
Starting the Administration Tool	13
Logging In	14
Administration Tool Interface	15
Menus	15
Main Information Window	17
Views	19
Settings	19
Viewing Loading Progress	22
Changing Column Display	23
Filtering Table Views	24
Inheriting parameter values	25
Part II. Nodes View	
5. The Circle	27
Node list	27
Node monitoring	28
Sampling period	30
Monitoring Options	30
Circle Properties	31
Editing Properties	33
6. Domains	35
Creating a New Domain	35
Domain Information	36
7. Nodes	39
Adding a New Node	39
Stopping a Node	42
Restarting a Node	42
Removing a Node	43
General Node Parameters	44
Cache	45
Cache Parameters	45
Clearing the Cache	46
Create a Consolidation Node	46
Sessions	47
Share association	48

Storage and Tape Pools	49
Policy	50
Network Interfaces	51
Properties	51
Logs	51
Monitoring	55
Service Management (Server Only)	55
Part III. Storage View	
8. Introducing Storage Pools and Devices	57
Storage Pools	58
Storage Devices	60
9. Local Disk Pools	63
Creating a Local Disk Pool	63
Pool Information	66
Disk Pool States	67
Disk Pool Availability	67
Pool Status Indicators	68
Cache Pools	68
Expanding Pools with New Partitions	69
Partition Information	69
Removing or Deactivating a Partition	70
Node and Policy Information	71
10. Shared Disk Pools	73
Creating a Shared Disk Pool	73
Pool Information	77
Changing Pool State and Availability	78
Pool Status Indicators	78
Adding Partitions to a Pool	79
Shared Partition Status	79
Partition Access Details	81
Deactivating or Removing a Partition	81
Managing Nodes in the Shared Disk Pool	84
Node Status Indicators	85
Changing the Order of Preference	86
Disabling the Order of Preference for Failback	87
Adding and Removing Nodes	89
Resetting the Pool	89
Validating the Pool Configuration	89
11. Tape Pools	91
Creating and Activating a Tape Pool	92
Activating a Tape Pool	94
Default Tape Pools	95
Unregistered Tapes	95
Blank Tapes	95
Cleaning Tapes	96
Excluded Tapes	96
Tape Pool Parameters	96
Tape Parameters	100
Advanced	102

Managing Tapes in Tape Pools	104
Search for Tapes	104
Tape menus and buttons	105
Assigning Tapes to Pools	107
Deleting or Removing Tapes	108
Checking Location Information	109
Moving Tapes to Mailboxes	109
Closing a tape	109
Integrity Check	109
Unused Tapes Check	111
Defragmenting and Duplicating Tapes	111
Importing an Archive	112
Export Pool	114
12. Cloud Pools	117
Creating a Cloud Pool	117
Cloud Pool Information	119
Cloud Pool Password	120
Cloud Pool States	120
Cloud Pool Availability	121
Cloud Pool Status Indicators	122
Node and Policy Information	122
13. Disk Devices	123
Partition Types	123
Adding a Local Partition	124
Add a Partition in File System mode	124
Add a Partition in Raw mode	126
Adding a Shared Partition	126
Shared Partition Groups	128
Deleting partitions	129
Create Pool from Partition	129
14. Library Devices	131
Library Detection and Information	131
Library Operations	133
Synchronizing a Library	134
Reinitializing a Library	134
Moving Tapes	134
Tape Drive Operations	136
Place Under Maintenance	136
Deleting a Tape Drive	137
Cleaning a Tape Drive	137
Access Nodes	138
Advanced Tape Drive Information	139
Part IV. Policies View	
15. Storage Policy	141
Create a Storage Policy	142
Configuring Storage Constraints	143
Defining a Rule	144
Shares Using the Policy	150
Replication Schedule	150

16. Versioning Policy	153
Default Policy	154
Create a Versioning Policy	154
Adding and Modifying Rules	156
Adding Rules	156
Understanding the Rules	158
Deleted Files Management	160
17. Archiving Policy	161
Archiving Properties	161
Create an Archiving Policy	162
Configuring Archiving Constraints	164
Policy Parameters	165
General	165
Filter	166
Advanced	168
Part V. Users View	
18. User Profiles and Access	171
Profiles	171
Directories and Protocols	172
Authentication	173
Access Control Overview	174
Access Rights on Shares	174
Unix Permissions	174
Access Control Lists (ACL)	175
19. Directory	177
Synchronization	177
Automatic Synchronization	178
Users	178
External Directories	179
Directory types and limitations	180
Order of Precedence	181
Merging Users and Groups	184
E-Mail Addresses	187
20. Importing External Directories	189
OpenLDAP Directory with POSIX Schema	191
Active Directory Imported in LDAP Mode	197
Active Directory Imported in Windows NT Mode	205
NIS Directory	208
21. Users and Groups	213
Viewing Users and Groups	213
Searching	213
Create a Local User	214
Modifying User Settings	216
Changing User Profile	217
Changing User Passwords	217
Managing Groups	218
User Tasks	218
Group Tasks	219
Share Access	220

User Sessions	222
Viewing Session Logs	224
E-Mail Addresses	225
22. Network Access Controls	227
Default Control	227
Create an IP Network Access Control	227
Defining IP Parameters	229
Create a Host Name Network Access Control	229
Apply a Network Access Control to a Share	230
Part VI. Shares View	
23. Shares	231
Getting Started	231
Create a Share	232
Create a Share Group	233
Move a Share	235
Rename a Share	235
Delete a Share	235
Policy and Archive Processing	236
Policy Status Indicators	237
Share Information	237
Modify Share Parameters	240
Policy Association	241
Policy Scheduling	242
Parameters	243
Access Controls	244
Activation on Nodes	247
Version Management	248
Quotas	250
Archives	253
Archive Details	255
Archive Display Filter	256
Manual Archiving	256
Connecting to Shares	256
Depositing or Migrating Files to Shares	258
24. Protocols	259
Activating the File Server	260
Restarting the File Server	261
Configuring the File Server	262
Access Timeout	262
Statistics	262
CIFS	263
FTP	264
NFS	265
Activating Shares	268
25. Auditing and Statistics	269
Activation	269
Share Audit Activation	269
Server Statistics Activation	270
Log File Characteristics	270

Location	271
Size limits	271
Duration and logging interval	272
Consolidation (audit only)	273
Log File Contents	273
Audit Log	273
Statistics Log	275
Part VII. Advanced View	
26. Distributed Lock Manager	277
How the Lock Manager Works	277
Lock Manager Administration	278
Nodes and Parameters	278
Viewing Lock Lists	280
Advanced	281
Conflict Resolution	282
27. Policy Management	283
Using Policy Management	283
Edit Parameters	283
28. HA Cluster Services	285
Election	285
Creating and Configuring HA Clusters	286
Adding Processing Nodes	287
General Settings	288
Cluster Types	288
IP Cluster	289
DNS Cluster	291
Scripted Cluster with OpenVPN	292
NetBIOS Cluster	293
Edit Parameters	293
Node Priority and Failback Mode	295
Network Interface	295
Network Settings	295
Share dependency	295
29. Outgoing SMTP Server	297
30. SNMP Target	299
Create SNMP Target	300
MIB	300
Part VIII. System Information	
31. Administration Log	303
Accessing the Administration Log	303
Log	303
Administration	308
32. Supervision	311
Note	312
Filtering the Notes	313
Supervision Note History	314
Working with Notes	314
Source Details	314
Note Registrations	315

Advanced	320
Share Note Count Limit	321
Part IX. Appendices	
Appendix A. Local Variables	323
Tape libraries	324
Network Interfaces	324
Encryption	325
Port Ranges	325
Multipath verification	326
Appendix B. Configuring and Testing Drives and Libraries	327
Configuring the SCSI Bus	327
Testing the Hardware Outside of Active Circle	327
Appendix C. Configuring external software for use with HA Clusters	329
Configuring DNS Clusters	329
Configuring Scripted Clusters with OpenVPN	329
Installing OpenVPN	330
Configuring OpenVPN	330
Running OpenVPN	330
Appendix D. Managing Windows ACLs	331
Editing ACLs	331
Appendix E. Migrating Files to Shares	333
Tools	333
Robocopy	333
Xcopy	333
Transferring files without Windows ACLs	334
Command syntax (without ACLs)	334
Synchronizing a Folder Structure	334
Transferring files with Windows ACLs	335
Prerequisites	335
Command syntax (with ACLs)	336
Xcopy syntax	337
Circle Properties	339
Index	341

Introduction

About this Guide

The purpose of this document is to provide system administrators with a comprehensive guide to the administration and configuration of the Active Circle storage system. The main focus is on the usage of the Active Circle Administration Tool application. The organization of this document reflects the layout of the Administration Tool interface.

Audience

This guide is written for Active Circle system administrators who want to understand the system architecture and components, and how to use the administration interface. To use this guide, you should have a system administration background.

Other documentation

The documentation for the Active Circle Storage System consists of the following documents:

- *Active Circle Administration Guide*
- *Active Circle File Explorer Guide*
- *Active Circle Installation Guide*
- *Active Circle Command Line Guide*
- *Active Circle Technical Overview*
- *Active Circle Technical Requirements*
- *Network port and routing requirements* (technical brief)
- Documentation for the Active Media Connector (AMC, the Active Circle API)
- Documentation for the Active Media Explorer (AME)
- Documentation for the Active Data Mover (ADM)



Note

Configurations shown in the application screens in this manual are for demonstration purposes only, and are not necessarily valid for real-life use. The sizes and capacities of disks, partitions, tapes, libraries and pools would be significantly greater in actual client installations. For more info on the requirements for Active Circle Storage System configurations, please refer to the *Active Circle Technical Requirements* document or the *Active Circle Installation Guide*.

Chapter 1. What's New

This chapter contains an overview of the most important changes related to the administration of the Active Circle Storage System from version 4.1 to version 4.5. For more detailed information about updates and bug fixes between versions, please refer to the *Release Notes*.

New and changed in version 4.5

The main changes can be summarized as follows.

- [Addition of Cloud Storage functionality](#)
- [Addition of LTFS archiving](#)
- [Updates to the Distributed Lock Manager](#)

You will find a short description of these product changes below.

Hybrid Cloud

Active Circle has introduced a hybrid cloud solution, allowing data to be archived to a cloud provider directly from the Active Circle Storage System. This has been implemented through a new type of storage pool called *Cloud Pools*. For more information, please see [Chapter 12, Cloud Pools](#).

LTFS Support for tapes

Active Circle now supports both reading and archiving using the LTFS tape format. LTFS can now be used for regular archiving in the same way as the Posix TAR format. A separate tape pool (or several) must be set up to be able to use the LTFS format for regular archiving. For more information, please see [Chapter 11, Tape Pools](#).

DRLM updates for problem resolution

Along with continued internal improvements in the distributed lock manager functionality, it is now also possible to stop (and restart) the DRLM services on a node without having to stop the node itself. This way conflicts and inconsistent locks can be resolved while the nodes of the Active Circle storage system continue running. For more information, please see [“Conflict Resolution”](#).

Chapter 2. Features and Concepts

Active Circle is a software platform designed to store, protect, and manage large and rapidly growing volumes of digital content. The Active Circle software platform combines different devices into a single and scalable pool of storage. It delivers a comprehensive set of integrated services for storage management, data protection, and data lifecycle management.

Active Circle has been designed for organizations that manage large volumes of data — video content, images, scientific and technical data — to help them better control the lifecycle of their data while at the same time keeping down the costs of storage ownership and administration.

Main Features

The main features of the Active Circle Storage System include:

- Hardware-independent software platform running on standard x86 servers.
- Data storage on disk or tape.
- Data storage in the cloud.
- Data access through standard network protocols — NFS, CIFS, and FTP
- An extensive application programming interface (API).
- A command line interface for simplifying and scripting repetitive tasks.
- Optimization through the use of hierarchical storage, hybrid cloud and data lifecycle management.
- Data protection and high availability provided by versioning, replication, and clustering.
- Multi-site data sharing.

Hardware-Independent Storage

The Active Circle platform aggregates servers and storage devices into an flexible and secure environment. Because it can be installed on standard x86 servers and supports any disk or tape storage system, Active Circle is modular and can be easily extended with the addition of storage space or servers. Your system can therefore evolve independently of storage technologies.

Data Protection By Design

The Active Circle system has built-in mechanisms for data protection, high availability, and business continuity planning. Data protection is ensured by file versioning, which allows users to retrieve previous versions in the event of an error. High availability is delivered by the Virtual File Server cluster: if one node fails, another automatically takes its place. Finally, business continuity is ensured through the replication of data over multiple remote sites.

Data Lifecycle Management

Active Circle automatically manages data according to rules defined by the system administrator. Depending on the options employed, Active Circle will store your digital content on the appropriate media according to your needs in terms of performance and capacity. You can store your most recent and most frequently accessed data on fast disks close to your users. Less urgent information can be placed on cheaper high-capacity disks. Lastly, tapes can be used for long-term, low-cost storage.

Storage Supervision

The Active Circle Administration Tool provides you with a single utility for supervising your storage network. The administrator can oversee the data volumes, available storage capacity, and rates of growth through the graphical interface and then allocate resources as required.

Key Concepts

Before proceeding to the rest of this manual, we recommend that you read the following sections. They provide an overview that will explain the terms used and the operations described in this document.

Storage Network

The Active Circle system is based on a network of storage nodes organized into three levels:

- The Circle
- Domains
- Nodes

The Circle

At the highest level is the Circle, the virtualized storage space created by Active Circle. It consists of all the servers on which the Active Circle software is installed. These servers, or Nodes (see below), are connected to one another in a network via TCP/IP and can be distributed over multiple physical sites. Multiple Circles can co-exist on the same network, each creating an independent and secure storage space.

Domains

A Circle is sub-divided into one or more Domains. A Domain is a group of Nodes linked either geographically or functionally that can be administered independently of other Domains. The Domains can communicate with each other to share resources, in particular in response to the rules defined in the Policies.

Nodes

A node is a server running the Active Circle software that provides three basic services:

- It functions as a storage resource in the Circle
- It provides a point of access to the Circle
- It performs task processing for the Circle

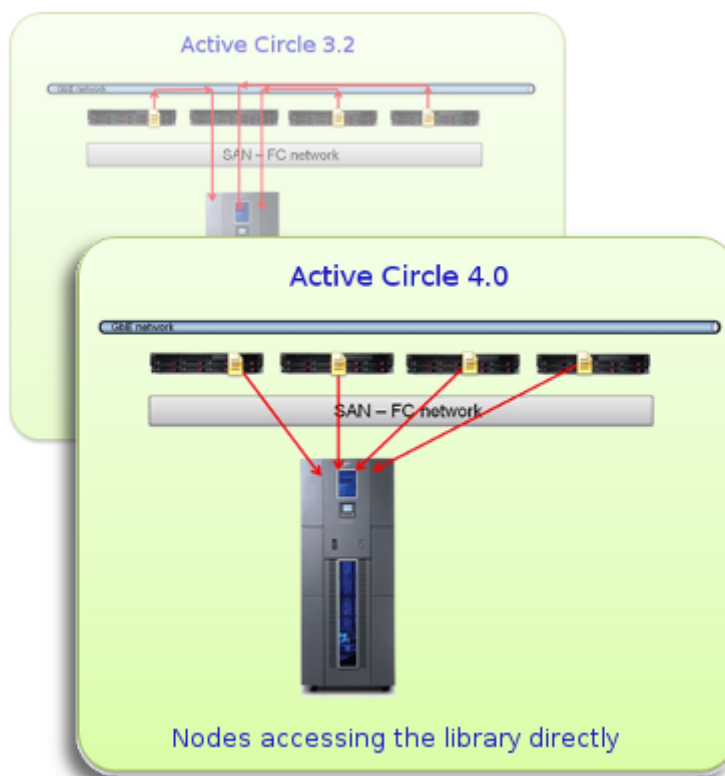
Nodes are logically grouped into Domains, and they can have various configurations in terms of hardware, software, storage peripherals, and operating systems. They can be located either on the same site or on multiple remote sites, sharing tasks according to the availability of their resources.

The functionality of the Active Circle system is distributed across all the Nodes and no individual Node is dedicated to a particular role. Due to this, the system is capable of resisting almost any kind of incident. For example, if a Node is removed from the Circle, its operations are automatically taken over by the other Nodes.

Consequently, end users work in virtualized storage spaces that guarantee the data security and access confidentiality.

Shared Tape Storage

In a SAN configuration, all the attached LTO drives will be available through direct access from all the Active Circle nodes. Drives will be dynamically allocated to nodes with data. This results in improved bandwidth and speed, since the data no longer need to pass through the node where the drives are attached. To make this possible, the Distributed Lock Manager and resource management have received new functionality, although this is not visible in the user interface.



Each node has direct access to tape libraries attached to a SAN.

Shared Disk Storage

In a SAN configuration with attached disk bays, the storage can be shared in failover/failback mode. Access is managed through an Active Circle node, but if this node becomes unavailable, the storage will automatically be accessible via another node in a process transparent to the users. This improves continuity of service in terms of storage access without the need for replication. For the administrator, this functionality introduces these features:

- Shared Disk Pools (see [Chapter 10, Shared Disk Pools](#))
- Shared Partitions (see [“Partition Types”](#))
- The concept of the Leader Node (see [“Managing Nodes in the Shared Disk Pool”](#))

Cloud Storage

Active Circle can be implemented as a Hybrid Cloud solution, with a first level of storage in-house and the second level hosted in a data center. This ensures data security and access both locally and in the cloud, with simplified management. The hybrid cloud functionality is based on the same technology and storage concepts as the system in general, and basically

consists of including standard Active Circle storage elements in Cloud storage pools, just as with disk and tape pools. For more information about defining Cloud storage, see [Chapter 12, Cloud Pools](#).

Unlimited Storage Pools

A storage pool is a space consisting of logical partitions mounted on any type of hardware. This space is subdivided into containers to store the data in an optimized form known as Packs. A pool groups together storage media of the same characteristics – SCSI, SATA, magnetic tape – in order to maintain the coherence of the data management rules defined by the policies. Each node can contain one or more data pools.

Storage resources can easily be increased by adding new storage peripherals to the existing nodes or new nodes to the circle. The new resources are immediately recognized without interrupting the service. The local archiving space can also be extended by using a hybrid cloud solution.

Integrity Check

Data integrity is systematically checked throughout the lifecycle of the files. A checksum based on the contents of the each file is generated when it is entered into the system and stored as metadata. This checksum is used to verify the integrity of the file each time it is read.

Data Transfer Optimization

In order to optimize data transfer over the network and reduce the storage volume of the Catalog, the metadata is compressed. This significantly speeds up data transfer rates during synchronization of the Nodes.

Policies

A Policy is a group of rules that are applied to a particular data storage location. Defined by the Active Circle system administrator, these rules determine:

- The number of copies of a file stored in the Circle
- The location of these files in the Circle
- The type of storage medium (disk or tape)
- The data management rules (how long to store files, whether and when to remove older versions, etc.)

Using the Policies, it is possible to define the lifecycle of your data according to the nature and criticality of the information. These rules can be modified as needed and the system responds immediately to the modification.

Integration

The Active Circle clustered file storage system has been designed to be integrated into the heart of your enterprise. The configuration information provided in this manual will allow you to adapt the storage network to meet your needs:

- Defining the enterprise domains
- Selection of functionality to be used: data replication, file versioning, and/or archiving.
- Combination of the Policies associated with this functionality.
- Implementing such additional features as high availability, etc.

Chapter 2. Features and Concepts

This manual describes each of these operations, and how to access the various resources required: networks, peripherals, directories, etc.

Chapter 3. The Active Circle Service

Starting the Active Circle Service

After installation and a reboot, the Active Circle service should start automatically. If the service is not running, you can start the service manually using one of the commands described below.

- In Red Hat or CentOS environments, you can execute the command:
`service activecircle start`
- In any GNU/Linux environment, you can execute the command:
`/etc/init.d/activecircle start`

If the service is already running, you can *restart* it by replacing the `start` option with `restart` in the commands above.



Important

You must start the service as `root`. If you do not, certain features will not be available, for example shares and HA Cluster Services.



Tip

If the Active Circle service does not start, it could be caused by one of the following situations.

- The date or time of the server is out of sync with the other nodes in the circle. Check the NTP configuration on the server. For more information, see the *Active Circle Installation Guide*.
- The server may be out of disk space for the Active Circle metadata.
- The node cache is full or unavailable.
- There are no IP addresses associated with the node's network interfaces.

Stopping the Active Circle Service

To stop the service, use the command that corresponds to your environment:

- In Red Hat or CentOS environments, you can execute the command:
`service activecircle stop`
- In any GNU/Linux environment, you can execute the command:
`/etc/init.d/activecircle stop`

Active Circle Service Extension

Service extensions allow system administrators to configure a sequence of scripts to be executed immediately after the start-up of a node and/or immediately before shutting it down.

This adds possibilities for users of the Active Circle storage system beyond basic mounting or unmounting shares. For example, it can be used at the start-up of the node to run preliminary

tests on the state of shares to perform network tasks such as firewall, routing, or multicast configuration; the availability of hardware such as tape libraries; or to perform statistical analysis on Circle activity. At shutdown, it can be used for checking jobs, checking the number of exceptions, warnings, and errors, to name just a few possibilities.



Important

This system should *not be used* for running interactive commands requiring user input because the output is sent directly to a log file.

When using this feature to run Active Circle commands, it is important to have previously run **acconfig** in order to configure the user login and password information and defined the default node and share against which to run the commands. Otherwise, the service start-up will wait for user input. For more information about the commands, please refer to the *Active Circle Command Line Guide*.

Specifications

The system works by executing the files located in particular directories. They are:

- [INSTALL_DIR]/service.d/start
- [INSTALL_DIR]/service.d/stop

The executable files located in the `start` directory are executed *AFTER* the start-up of the Active Circle. The executable files located in the `stop` directory are executed *BEFORE* the shutdown of the node.



Note

Running the command **service activecircle start** (or its equivalent) on a node that is already running will not trigger the execution of the files in the `start` directory. Similarly, running the command **service activecircle stop** (or its equivalent) on a node that is stopped will not trigger the execution of the files in the `stop` directory.

There are some constraints on which files are executed and how:

- The system supports any type of executable: script, command line, C program, Java program, etc.
- Only executable files are taken into account by the system. Clients can therefore store older versions of their scripts in sub-directories, if they so choose. Any executables found in sub-directories will not be executed.
- Executable files *must not* have spaces in their names.
- Files are executed sequentially according to their alphanumeric order in the directory.
- Files with names ending in a tilde "~" are ignored.
- For each file executed a line is displayed with the status of the operation (success or failure) depending on the return code.
- If one of the executables terminates with a return code other than 0 (zero), all subsequent commands are ignored and one of the following messages is displayed:

```
Error executing service extension startup scripts (subsequent
scripts won't be executed).
```

Error executing service extension shutdown scripts (subsequent scripts won't be executed).

This does not have an effect on the start-up or shutdown of the node.

Implementation

Perform these steps to implement the service extensions:

1. Create the directories for the executables:
 - `[INSTALL_DIR]/service.d/start`
 - `[INSTALL_DIR]/service.d/stop`
 - `[INSTALL_DIR]/service.d/logs`
2. Copy the executable files to the appropriate directories. Make sure that the files are executable, that they are in the proper (alphabetic) order of execution, and that their names respect the constraints listed above.

The next time you use the command **service activecircle start | stop | restart**, the system will launch the scripts accordingly.

The log files will be written to the directory `[INSTALL_DIR]/service.d/logs` and the name of each log file will begin with the run date.

Chapter 4. The Administration Tool

The Administration Tool is the application interface for configuring and administering the Active Circle storage system. You can use it to configure:

- The Circle, its Domains, Nodes, Pools, and Partitions
- The policies for data management, file versioning, and archiving.
- The access to tape resources, including SCSI bus detection by Active Circle.
- The session and network access control using local and/or external directory services.

Starting the Administration Tool

You can access the Administration Tool by launching the client in one of two ways: either locally on one of the nodes or remotely via a web browser with the Java Web Start plug-in.

To run the Administration Tool directly on a node, execute the following command:

```
/activecircle/admin/bin/admin &
```

To start the Administration Tool from a web browser, type the server name or IP address of an Active Circle node into the address line of the browser and press **Enter**. The Active Circle main page appears with links to the Administration Tool and the File Explorer:

Figure 4.1. Active Circle Main Page



Click on the **Administration** link to start the Administration Tool application.



Note

To run the Administration Tool via a web browser, you must have the Java Web Start plug-in installed. If you do not have the plug-in, you can download it by clicking on the Java™ icon at the bottom of the page.



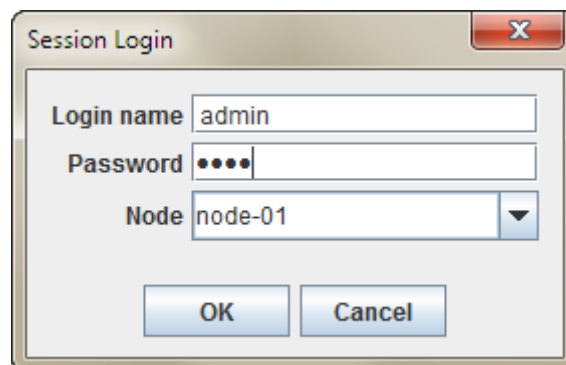
Tip

If you are connecting from a Microsoft Windows machine that is running the Windows firewall, be aware that this can sometimes interfere with the Administration Tool application. If this is the case, you should add the application to the list of exceptions in the firewall.

Logging In

Type the user name and password in the login dialog box. The default system administrator name is `admin` and the password is `1234`.

Figure 4.2. Login dialog box



The image shows a 'Session Login' dialog box with a title bar containing a close button (X). Inside the dialog, there are three input fields: 'Login name' with the text 'admin', 'Password' with four dots, and 'Node' with a dropdown menu showing 'node-01'. At the bottom, there are two buttons: 'OK' and 'Cancel'.

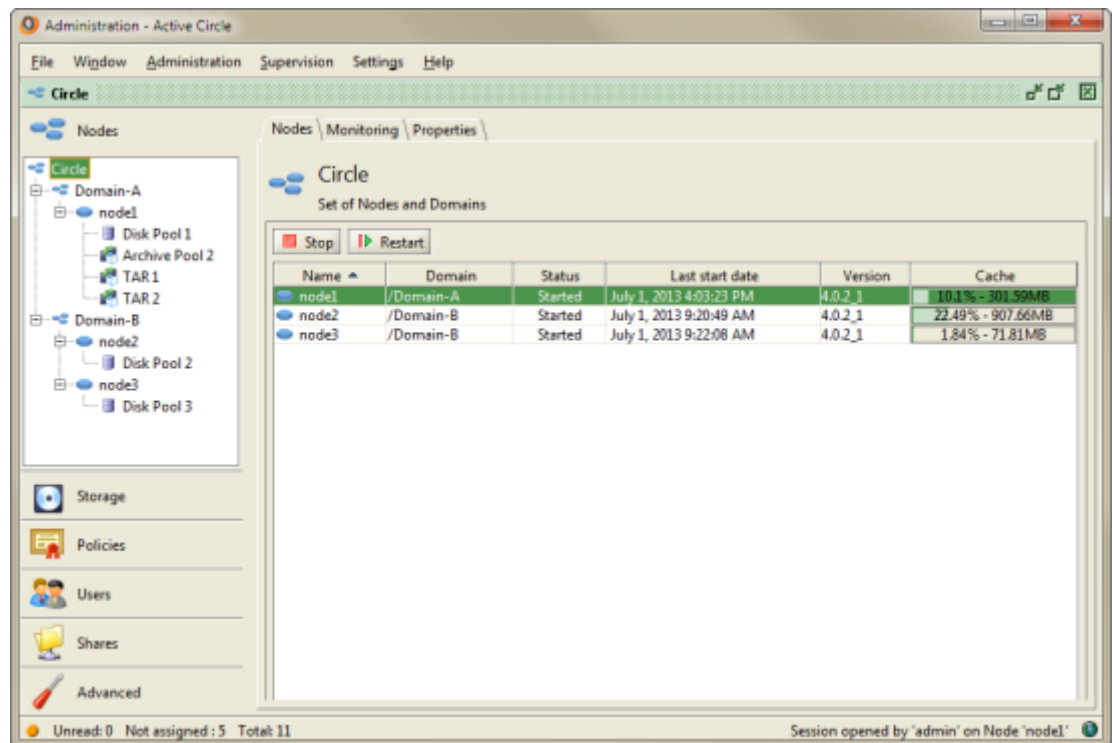


Change Password for the `admin` User

You should change the default password associated with the `admin` account immediately after installation. For information, please refer to [“Changing User Passwords”](#) in [Chapter 21, Users and Groups](#).

Administration Tool Interface

Figure 4.3. The Administration Tool Interface



When the Administration Tool opens, the user is presented with a standard application interface. There is a row of menus (see “[Menus](#)” below) across the top of the window. Under that is the Main Information window (see “[Main Information Window](#)” below).

Menus

This section describes the menus in the Administration Tool.

Table 4.1. File Menu

Menu option	Usage
Login	Connect to the system. You will be prompted for your username and password. Available only if you are logged out.
Logout	Logs you out of the system. Available only if you are logged in.
Exit (Ctrl+Q)	Logs you out of the system and closes the Administration Tool.

Figure 4.4. The Window Menu

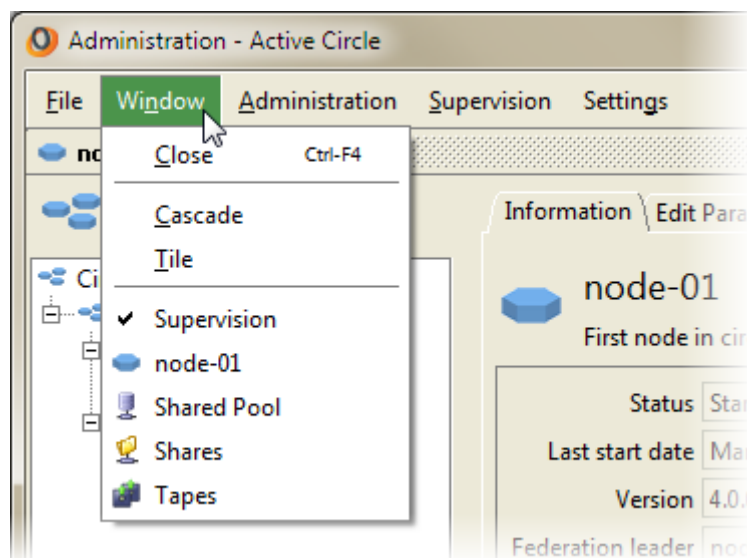


Table 4.2. Window menu

Menu option	Usage
Close (Ctrl+F4)	Closes the current window.
Cascade	Arranges all open windows in an overlapping style.
Tile	Arranges all open windows side-by-side.
<i>(Window list)</i>	All open windows are listed at the bottom of the menu. Select a window from the list to bring it to the foreground.
Full Screen	Maximizes the application window so that it covers the whole computer screen. If you want to cancel the full screen mode, select the menu option again to remove the check mark.

Table 4.3. Administration menu

Menu option	Usage
New Window (Ctrl+N)	Opens a new instance of the Main Information window. It is possible to have multiple information windows open within the application at the same time.
Administration Log (Ctrl+Shift+L)	Opens the Administration Log window. The Administration Log is the interface for viewing and managing the system's logging services (see Chapter 31, Administration Log for more information).

Table 4.4. Supervision menu

Menu option	Usage
New Window (Ctrl+Shift+N)	Opens a new instance of the Supervision window for viewing Supervision Notes (for more information, see Chapter 32, Supervision).

Table 4.5. Settings menu

Menu option	Usage
Save Settings (Ctrl+S)	Saves the information about the current arrangement of the various windows.
Export...	This option launches a wizard that allows you to export all of the information about the configuration of your Circle to an XML file.
Preferences	This menu item opens a dialog box that allows you to set your preferences about the behavior of the Administration Tool interface and its look and feel. For more information, see “Preferences” .

The contents of the **Settings** menu is covered in more detail in [“Settings”](#) below.

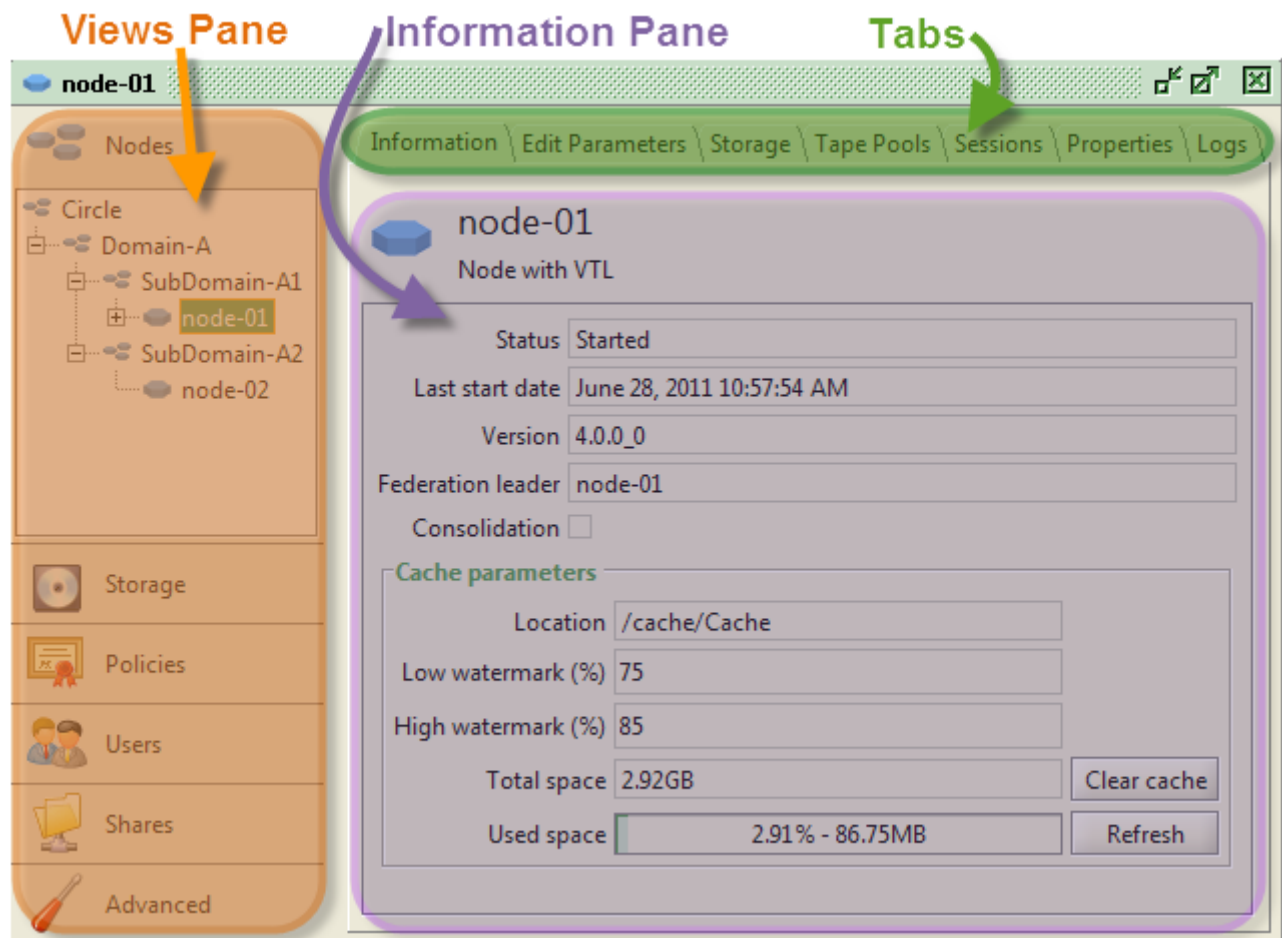
Table 4.6. Help menu

Menu option	Usage
Documentation	Opens a browser window giving access to the Active Circle documentation.
License	Opens the License Information window, where you can check your Active Circle license data and update the license after receiving a new license key from Active Circle.
About	Shows the Active Circle software version and patch level.

Main Information Window

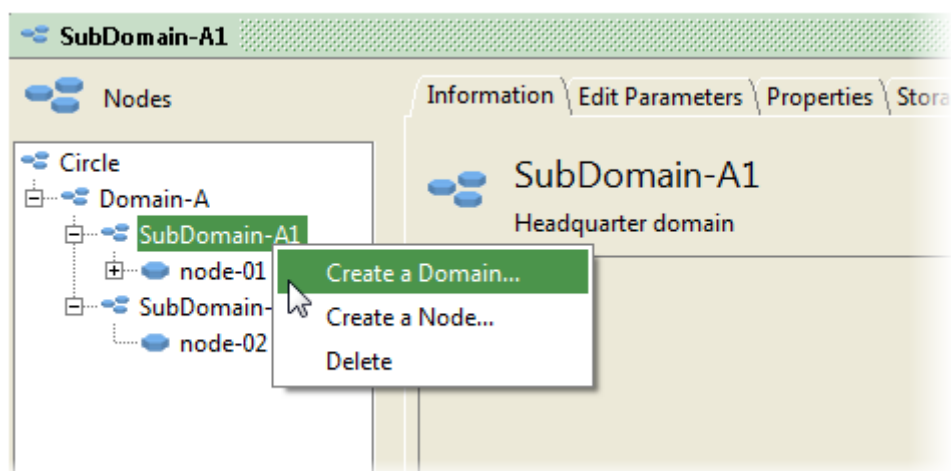
The Main Information window provides access to all of the functionality of the Active Circle Administration Tool. The window is divided into two panes: the left-hand pane presents a column of *views* (see [“Views”](#)) of the different categories of objects managed by the system. When you click on a view, it opens to reveal the elements relevant to that particular view. Clicking on an element within a view populates the right-hand pane with information relevant to that particular element presented in a series of *tabs*. These tabs can be used to display the information and change the configuration for each element.

Figure 4.5. Administration Tool - Panes and Tabs



In addition to the tabs, some elements have a *context menu* which can be displayed by right-clicking on the element itself in the views pane. The context menu can be used to perform actions on the elements, typically actions related to creating or deleting elements.

Figure 4.6. Administration Tool - Context menus









You can manually adjust the size of the panes (and in some cases the zones inside a tab) using the mouse.

Views

The views provide access to most of the essential functionality in the Active Circle system. Most administrative tasks can be performed by selecting an element in the view pane, and then using the associated information pane and tabs to configure the settings.

Table 4.7. Administration Tool Views

View	Contents
Nodes 	This is the default view on opening the application. The Circle, its Domains and sub-Domains, the Nodes and their Storage Pools are presented in a tree structure. For more information, see Chapter 5, The Circle .
Storage 	This view lets you manage the storage pools, disks and tape libraries attached to your system. For more information, see Chapter 8, Introducing Storage Pools and Devices .
Policies 	Use this view to configure policies for the system. For more information, see: <ul style="list-style-type: none"> • Chapter 15, Storage Policy • Chapter 16, Versioning Policy • Chapter 17, Archiving Policy
Users 	This view is used for managing users and groups, and configuring network access controls. For more information, see: <ul style="list-style-type: none"> • Chapter 18, User Profiles and Access • Chapter 19, Directory • Chapter 21, Users and Groups • Chapter 22, Network Access Controls
Shares 	The Shares view; allows you to create and configure the Shares that are the points of connection for end-users to the system. For more information, see Chapter 23, Shares .
Advanced 	The Advanced view allows you to manage and configure the features that are described in these chapters: <ul style="list-style-type: none"> • Chapter 26, Distributed Lock Manager • Chapter 27, Policy Management • Chapter 28, HA Cluster Services • Chapter 29, Outgoing SMTP Server • Chapter 30, SNMP Target

Settings

The **Settings** menu contains options for configuring and managing the Administration Tool.

Save settings

The **Save settings** option allows you to save the current arrangement of the application window so that when you return to the system, regardless of which method or machine you use to connect, the application window will appear as it did when you saved it. The settings will be saved for the current user.

To save your current settings, click on **Settings** in the menu bar and select **Save Settings**. You can also use **Ctrl+s**.

Export

The **Export Wizard** allows you to export the information about the current configuration of your Active Circle environment to an external file. Although you can choose the format of the exported file - XML and/or binary - note that the binary option is reserved for future use; that is, there is no option for re-importing the binary file.



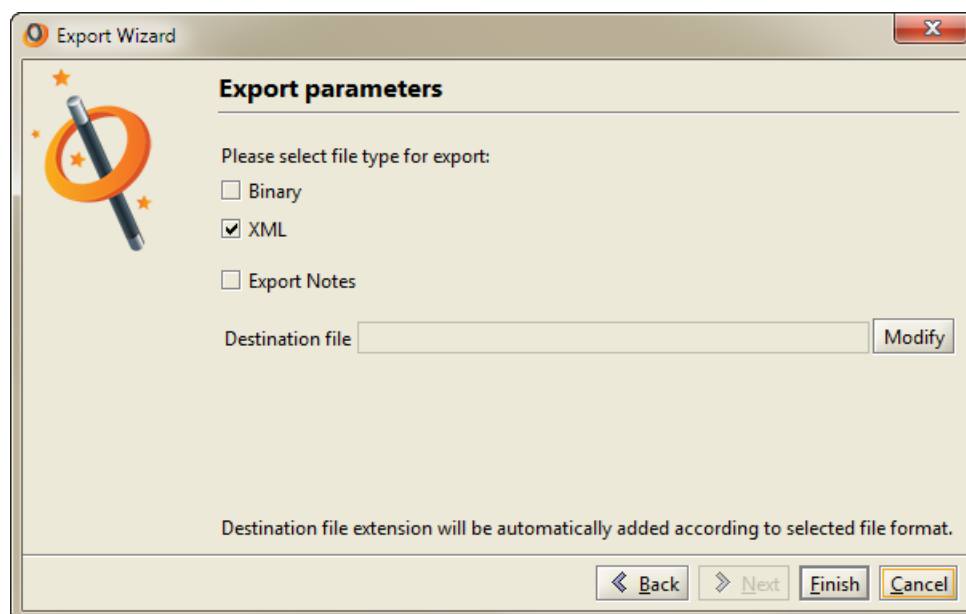
Tip

The information exported to the XML file can be helpful in the event of a troubleshooting operation and may be requested by an Active Circle support engineer.

To launch the **Export** wizard, click on **Settings** in the **menu bar** and select **Export....** This displays an interactive window that will lead you through the export process.

Click the **Next** button.

Figure 4.7. Export Wizard



The option **XML** should be selected by default. In the **Destination file** field, click the **Modify** to browse to the location where you want to save the file, and then enter a file name without an extension (the wizard will provide the appropriate extension according to the file type).

You can also choose whether or not to include the **Supervision Notes**. For more information about **Supervision**, see [Chapter 32, Supervision](#).

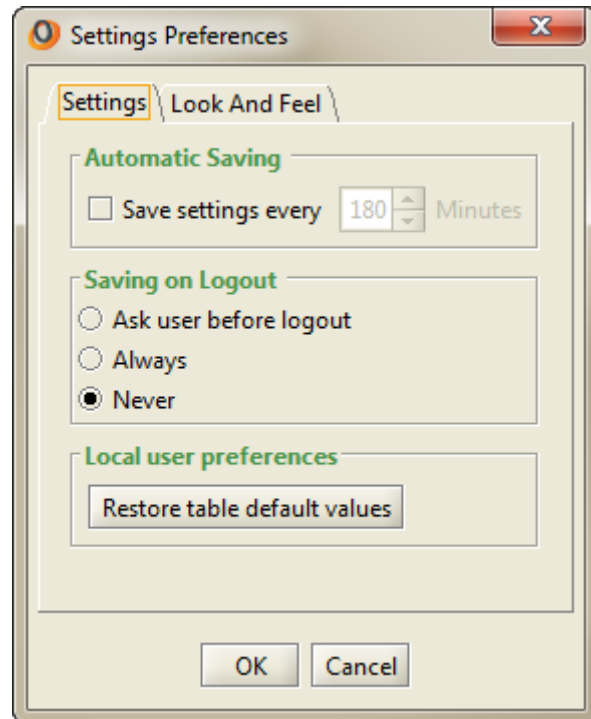
When you are ready, click **Finish**. The application does not inform you that the task has been completed. Go to the directory to which you exported and you can browse the XML content.

Preferences

The Preferences menu options allows you configure how the **Settings** are saved as well as the look and feel of the Administration Tool.

To access the Preferences dialog box, click on **Settings** in the **menu bar** and select **Preferences**.

Figure 4.8. Preferences: Settings tab

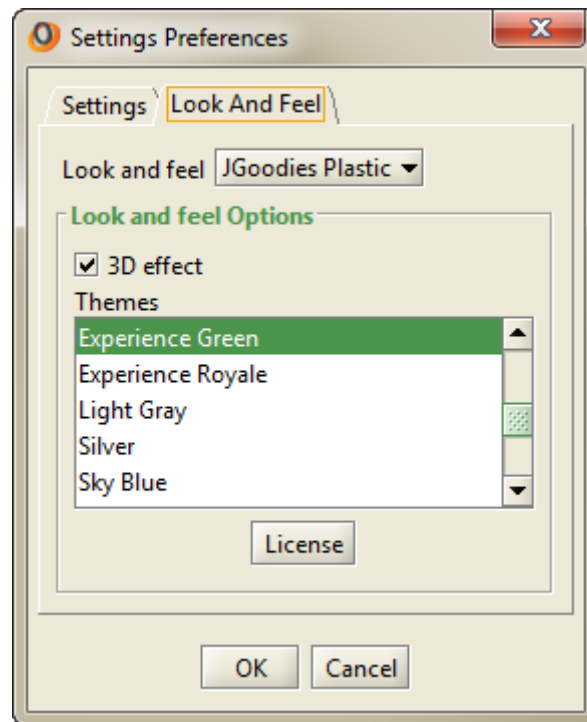


In the **Settings** tab, you can configure the application to save your settings automatically by clicking the **Save settings every...** check box and selecting a time (in minutes).

You can also configure the application to save your settings when you log out by choosing either **Ask user before logout** or **Always**.

Under **Local user preferences**, you can reset all table view back to the default settings by clicking on the Restore table default values button.

Figure 4.9. Preferences: Look and Feel tab

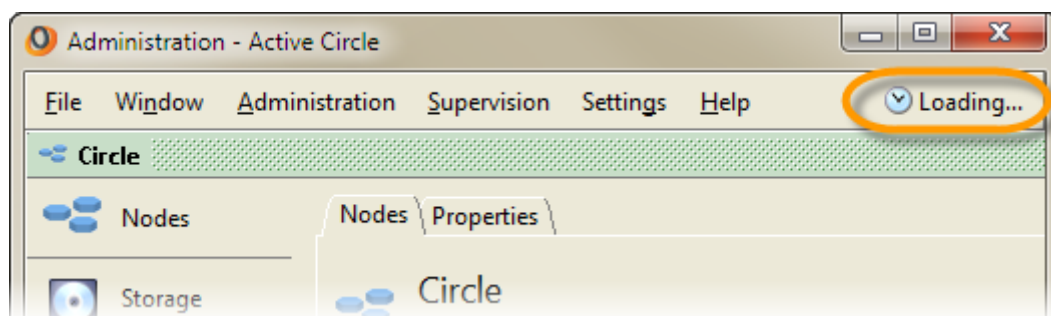


The **Look and Feel** tab allows you to choose from several different themes and skins to change the appearance of the user interface. The available choices depend on which operating system you use and the desktop appearance you have configured. You may have to restart the application for all the look and feel changes to take effect. The preferences will only be changed for the current user, so each user can select a different graphical theme.

Viewing Loading Progress

When the Administration Tool is launched, an indicator is shown in the top right corner of the application window to indicate that some elements of the graphical user interface are still being loaded:

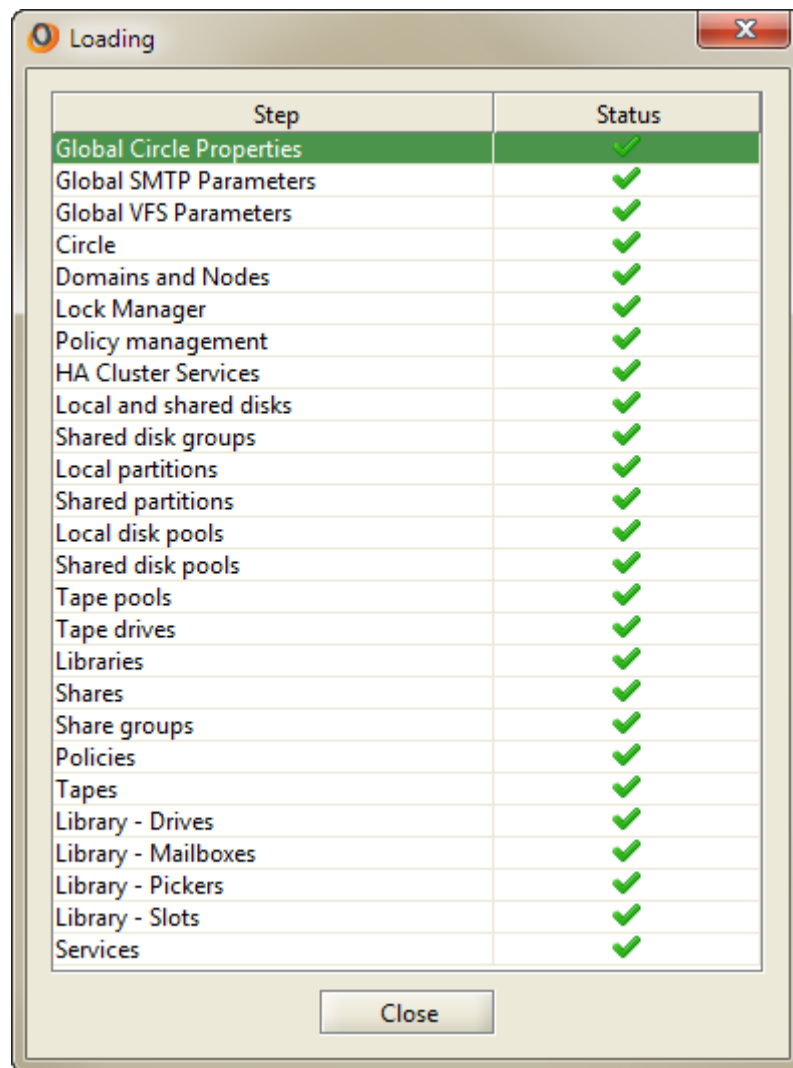
Figure 4.10. Loading Indicator



While the indicator is being shown, you can click on it and select **Details** to display a list of all the interface loading steps and their status.

When all the elements are loaded, this indicator disappears.

Figure 4.11. Loading Step List

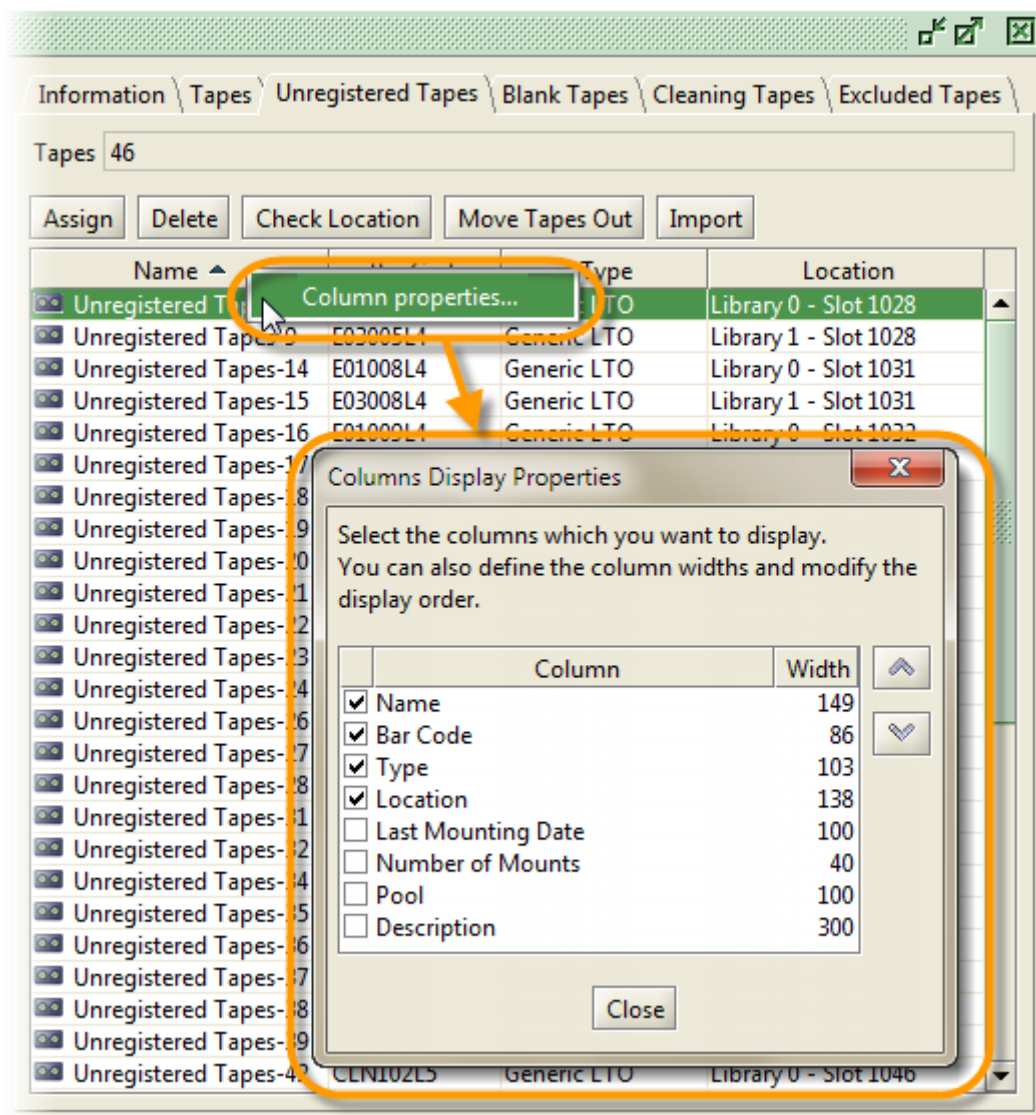


Step	Status
Global Circle Properties	✓
Global SMTP Parameters	✓
Global VFS Parameters	✓
Circle	✓
Domains and Nodes	✓
Lock Manager	✓
Policy management	✓
HA Cluster Services	✓
Local and shared disks	✓
Shared disk groups	✓
Local partitions	✓
Shared partitions	✓
Local disk pools	✓
Shared disk pools	✓
Tape pools	✓
Tape drives	✓
Libraries	✓
Shares	✓
Share groups	✓
Policies	✓
Tapes	✓
Library - Drives	✓
Library - Mailboxes	✓
Library - Pickers	✓
Library - Slots	✓
Services	✓

Changing Column Display

In many of the information panels in the Active Circle Administration Tool interface, the data are presented in table form. Each column has a heading, and clicking on the heading will sort the information based on that column. In most of these tables, you can also select which columns to display. You can add more information columns (if available for the particular element), remove data columns and change the column order and widths. To modify the columns display, right-click on any column heading and select **Column Properties** from the context menu. Then select which columns you want to have displayed from the list in the **Columns Display Properties** window. Click Close when you are done.

Figure 4.12. Changing Column Properties



The changes will be saved locally, so the column display will be the same each time you start the Administration Tool on the same client.

Filtering Table Views




Certain information tables have a *filtering* feature that allows you to filter the data displayed to make it easier to find the relevant information. Examples are the lists of users, groups and tapes, the node Session tab and the circle properties lists. The filter button  is located at the top right corner of the table, in the column header line. It will only be displayed if the table contains enough information to make vertical scroll bars appear.

Figure 4.13. Filter Button



When you first open a tab that has the filter button enabled, the arrow inside the button will be blinking green. Click on the button to activate the entry fields for filtering.

Figure 4.14. Filter Fields




The number of entry fields will depend on the number of columns in the table. Type the characters or numbers you are searching for, into one or more of the filtering fields. This will reduce the contents of the table to display only what you want to see. The table contents will be updated as you type.

To end filtering the contents of the table, click on the 'X' that appears under the filter button to the right of the entry fields.

Inheriting parameter values

The interface is divided into views, which contains several items and sub-items. Most of these items have settings and parameters that can be defined and modified by the administrator. For many of the fields where values for the settings are specified, there are additional check boxes to the right of the original value field without any descriptive text. These are *Inherit* check boxes. When one or more of these boxes are checked, it means the parameter will inherit the value for this field from its parent object (a group, a root object, etc.). You will find these boxes to the right of many fields throughout the interface, they are always used for inheriting values. Some are only visible when an object is part of a group or a domain where values can be inherited. Others are dependent on default values set by circle properties. When the *Inherit* box is checked, the original value field is greyed out.

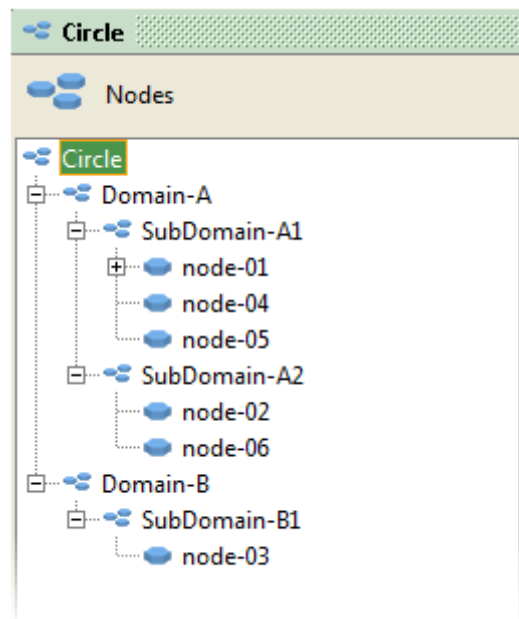
Figure 4.15. Example of Inherit check boxes



Chapter 5. The Circle

The Circle is the highest level object in Active Circle. All of the objects used by the system — Domains, Nodes, Pools, etc. — are located under the Circle. The **Nodes** view displays the organization of the Circle as a hierarchical tree structure. The circle will always contain at least one domain, where the nodes can be organized under additional sub-domains if needed. A domain can for example be a geographical site in a corporate structure. When configured, the Nodes view will show all storage pools in the Circle as elements under the appropriate node.

Figure 5.1. Circle elements

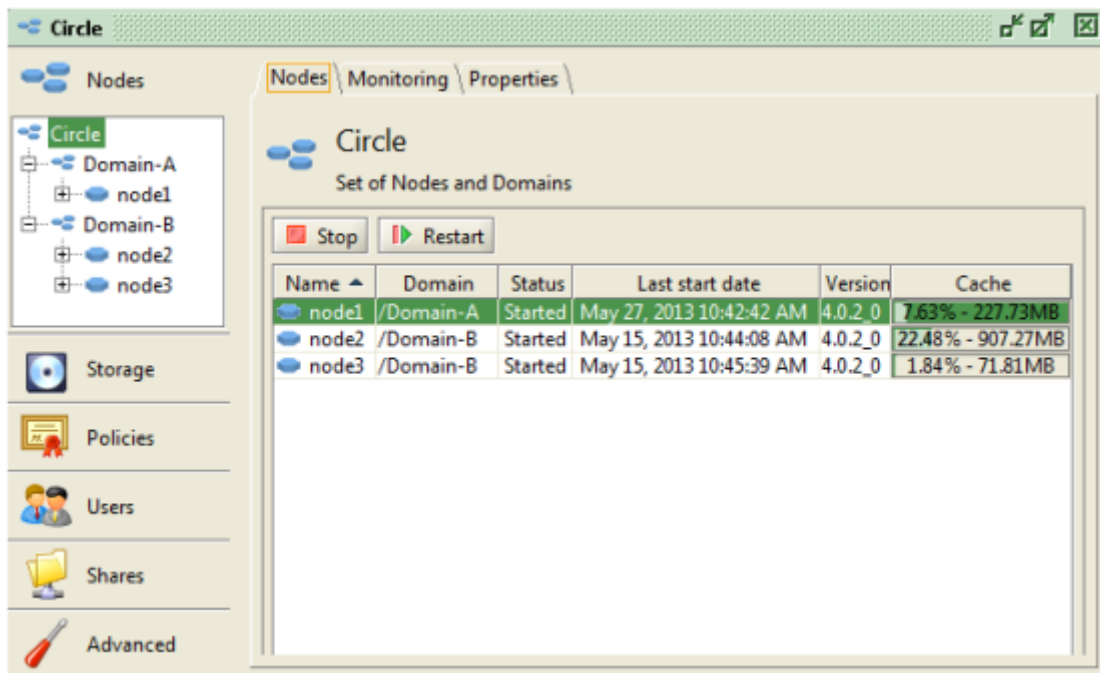


For most organizations, a single Circle is sufficient for handling their storage needs, although it is possible to implement multiple Circles within the same network. Note that these Circles will not be mutually aware. Any transfer of files from one to the other would occur as a regular network file transfer operation.

Node list

The **Nodes** tab of the **Circle** shows the list of all the nodes in the circle, with summary information. This gives the administrator a quick overview of all the nodes and a mechanism to stop and restart each one.

Figure 5.2. Overview of Nodes in the Circle



To restart a node, select it in the list and click on the **Restart** button. The Active Circle service will be restarted on the node.

To stop the Active Circle service on a node, select it in the list and click on the **Stop** button. The Active Circle service will be terminated on the node. Note that once a node has been stopped, you cannot start it again using the administration interface. To start it, you must use a terminal session on the server, as described in [“Starting the Active Circle Service”](#).

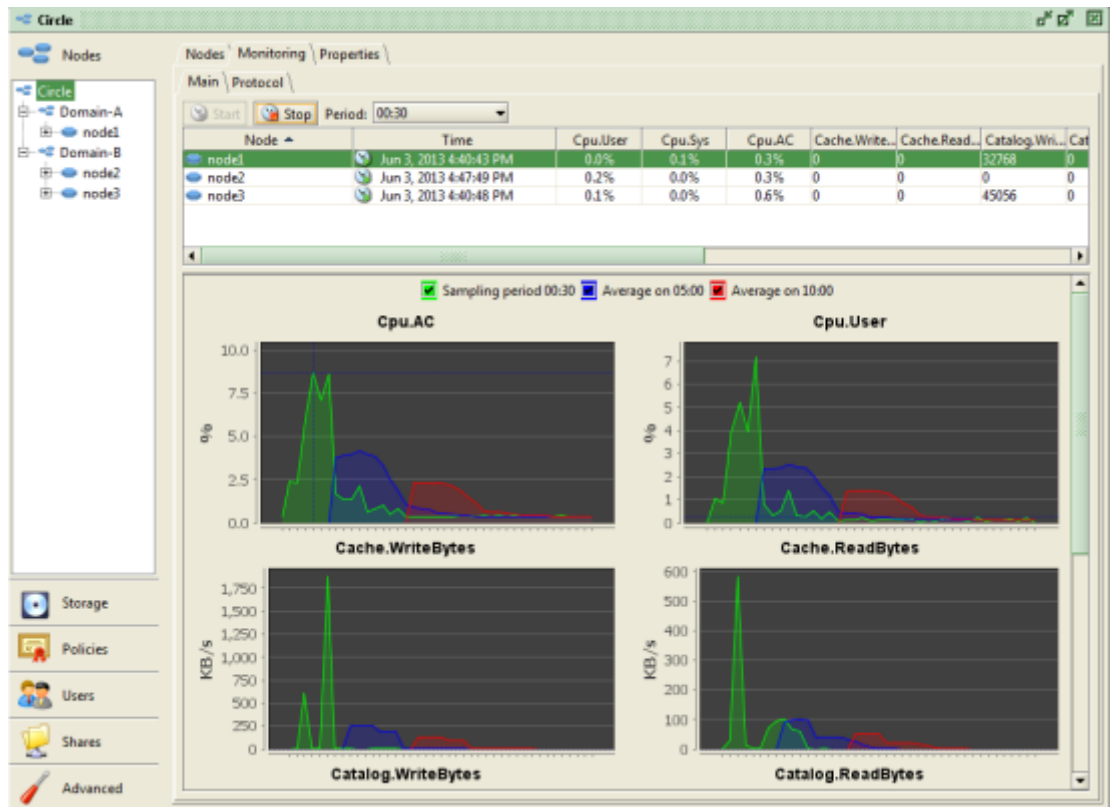
The data columns for each node contain the following information:

Node name	Last start date
Domain	Software version installed
Service status	Cache information

Node monitoring

You can monitor the activity and performance of the nodes in the circle by using the **Monitoring** tab. In the **Nodes** view, select the **Circle** item and go to the **Monitoring** tab. The same tab exists if you select a domain or a single node.

Figure 5.3. Monitoring tab



The monitoring interface consists of two tabs, **Main** and **Protocol**. Both tabs have the same features:


- Buttons for starting and stopping the monitoring, and a sampling period selector that is activated once the monitoring is underway
- A list of all the nodes in the circle, showing monitoring data in table format for all counters
- Graphical diagrams for certain performance counters. Each diagram has a right-click context menu with display and print options.

The **Main** tab contains performance counters for a node's CPU, cache, catalog, memory, swap file and load.

The **Protocol** tab contains counters for data read from and written to the node through the network. The information is presented as totals as well as broken down per protocol, requests and cache.

To see a descriptive text for each counter, place the mouse pointer on a column heading or on a diagram, and a tool tip will be displayed:

Figure 5.4. Counter description

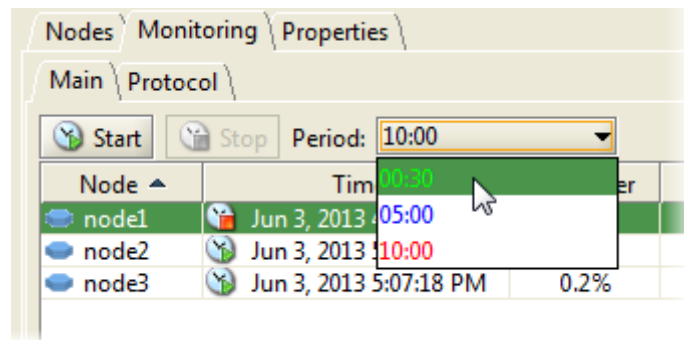


To start monitoring, select a node in the list and click on the **Start** button. The current time will be displayed in the **Time** column in increments defined by the sampling interval (30 seconds by default). To stop monitoring a node, select the node and click on the **Stop** button. The **Time** column will then show the time when the monitoring was stopped for the node.

Sampling period


Whether the monitoring is running or has been stopped, you can change the interval of the performance data being displayed in the table by using the **Period** drop-down box above the list of nodes. The default intervals are **30 seconds** (sampling interval), **5 minutes** (average) and **10 minutes** (average). When a period is selected, the data in the performance data table change to reflect the interval chosen. Note that the **Period** menu is only functional once the monitoring has started and the time of the first sampling period has passed.

Figure 5.5. Sampling period menu



The intervals and color coding shown in the **Period** menu correspond with the graphs being drawn in the charts in the lower part of the monitoring windows. Above the diagram area, you can select which graphs to display. You can choose to display one or several at the same time. Place a check mark in the box for the period you want to display. The first box, representing the basic sampling interval (green graph), is checked by default.

Figure 5.6. Sampling periods to display in diagrams



Monitoring Options

In addition to the sampling intervals described above, you can change the display of the data both in the table section and the diagram section of the monitoring tabs.


Column Display Properties

By default, all the available counters are displayed in the table part of the monitoring interface. You can choose which columns to display by right-clicking on any of the column headings and clicking on **Column properties**. The **Columns Display Properties** window will be shown, where you can select or deselect columns and define the default width of each column. Click on **Close** when you have made your selection.

Copy Chart

You can copy a chart from the monitoring tab to another application, for example a word processing document or a presentation program. Right-click on the diagram you want to copy, and select **Copy** from the context menu:

Figure 5.7. Copying a Chart



Save Chart

You can save a chart from the monitoring tab as an image file (PNG format). Right-click on the diagram you want to save, and select **Save as** from the context menu. Select a location and enter a name for the file.

Print Chart

You can print a chart from the monitoring tab. Right-click on the diagram you want to print, and select **Print** from the context menu.

Zooming and selecting the display range

You can zoom in and out on a diagram and select the range you want to display. You find these options by right-clicking on a diagram and selecting either **Zoom In**, **Zoom Out** or **Auto Range** from the context menu. If you select **Auto Range** and then **Range Axis**, the scale of the vertical axis will be reset so that the whole chart is displayed.

Instead of using the menu, you can also use your mouse to select an area of the chart to zoom in on. You can also use the scroll button on your mouse to increase or decrease the range of the vertical axis.

Circle Properties

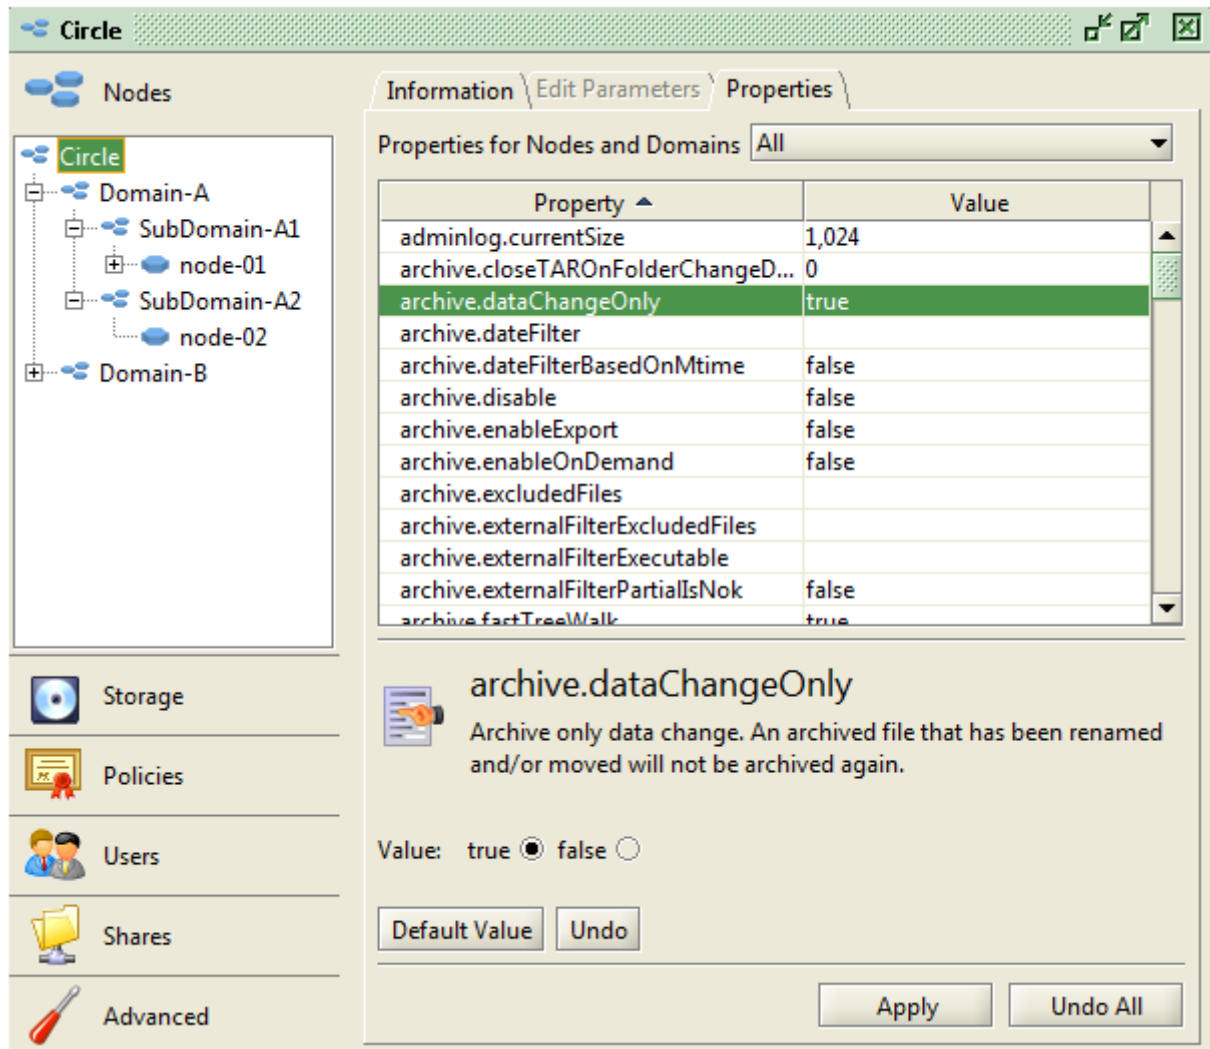
The Active Circle **Properties** enable you to get the best performance from your Active Circle configuration by allowing you to configure the behavior of the Active Circle nodes. This can be done at the level of:

- The Circle. In this case the configuration options will apply to all the Domains and Nodes in the tree structure.
- The Domain. In this case the configuration applies to all the Nodes within the Domain.

- The individual Node. In this case the configuration only applies to the Node.

To view the **Properties** for the Circle, click on the Circle element in the Nodes view and then click on the **Properties** tab.

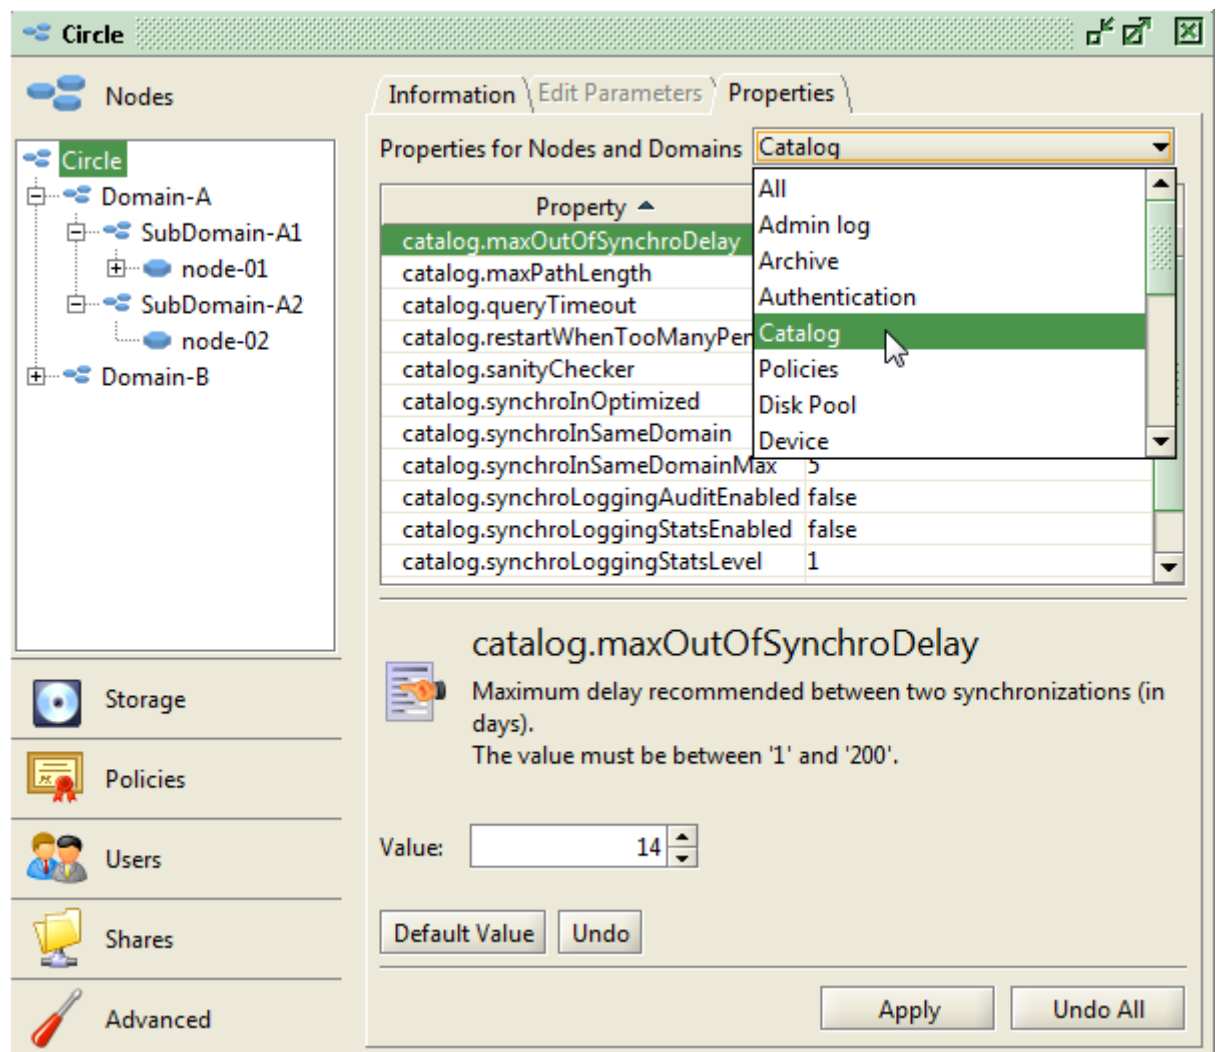
Figure 5.8. Circle Properties tab



You can filter the **Properties** view by clicking on the drop-down menu labelled **Properties for Nodes and Domains** and selecting one of the categories listed there. The possible categories are:

- | | |
|----------------|--|
| All | Distributed Resource Lock Manager (DRLM) |
| Admin log | HTTPD / Java Web Start |
| Archive | Integrity checker |
| Authentication | Lease |
| Catalog | NAS |
| Policies (COS) | Packs |
| Depository | Session |
| Device | Supervision |
| Disk Pool | Tapes (Media) |
| Directory | VFS |

Figure 5.9. Circle Properties tab: Filter



You can also filter the table by any character by using the [filter button](#) which appears in the upper right corner of the table.

The properties are by default inherited by all the domains and nodes within the circle.



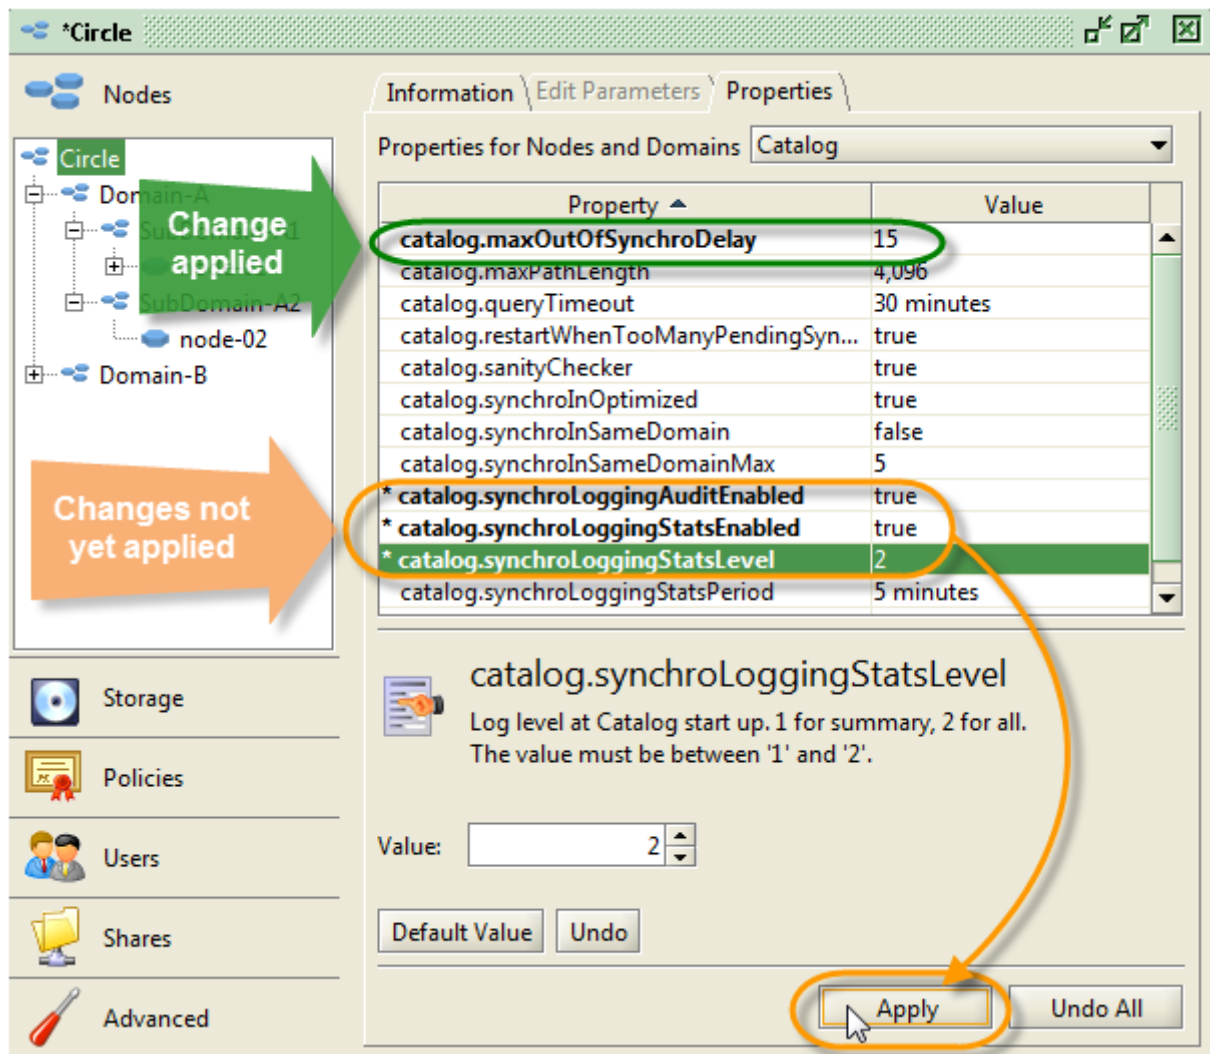
Tip

You will find a list of all the circle properties described in this manual in the index at the end of this document.

Editing Properties

You can edit the **Properties** by selecting one of the properties and then making the appropriate changes in the edit area at the bottom of the tab.

Figure 5.10. Circle Properties tab: Edit



When you make a change to a property, the modified line is displayed in bold. An asterisk appears at the beginning of the line to indicate that the property has been modified, but that the change has not yet been applied. Click the **Apply** button to validate the change. Click **Undo** to cancel the change. Clicking **Undo All** cancels all changes that have not yet been applied. You can always return a property to its default value by clicking the **Default Value** button.



Tip

You can also use the command `accircleproperty` on the Active Circle Command Line to manage circle properties. For more information about this, see the *Active Circle Command Line Guide*.

Chapter 6. Domains

The domains and sub-domains are displayed under the Circle in the Nodes View. The nodes and their associated storage pools are attached directly to the domains or sub-domains.

Domains help you to manage the Nodes in your storage cluster (or Circle) by allowing you to organize them into a hierarchical structure according to the needs of your organization. The Circle domains could for example be organized geographically.

A Circle must have at least one domain and the creation of one is required as part of the installation process. It is possible to create multiple domains and even to create sub-Domains. There is no limit to the number of Nodes a Domain may have. These subjects will be covered in more detail in the following sections.

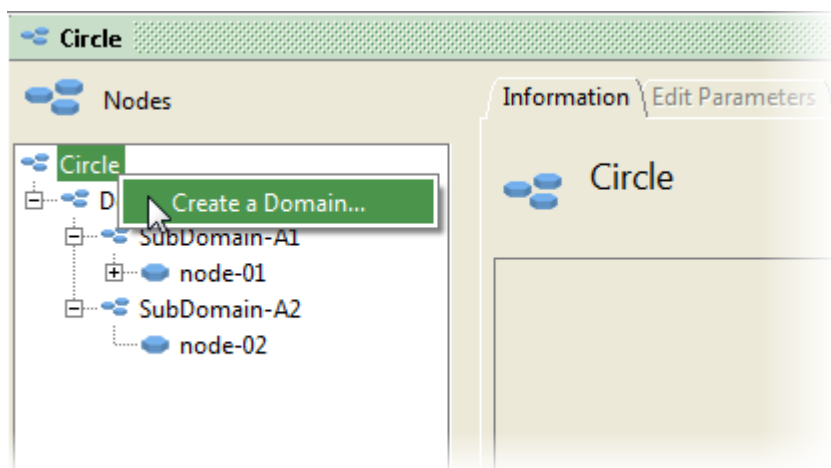
Creating a New Domain

Because a Domain is created during the installation procedure of the first Node, this step is optional. It is possible to locate all of your Nodes in a single Domain. However, if you plan to install a significant number of Active Circle Nodes and these Nodes are distributed over different geographic sites, it is probably a good idea to represent these sites as Domains. These additional Domains and/or sub-Domains can be created through the Administration Tool.


To create a new Domain, follow the procedure below.

Procedure 6.1. Create a New Domain

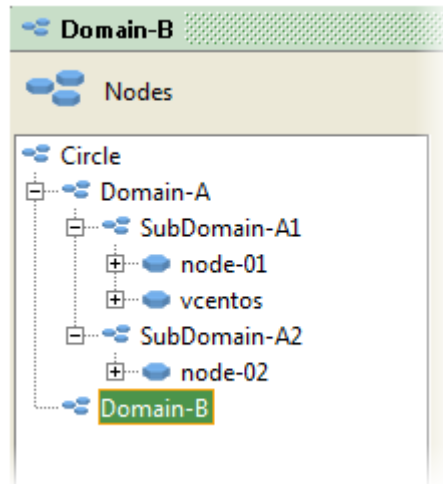
1. Right-click on either the Circle or one of the Domains and click on **Create a Domain....** Selecting the top-level object, Circle, will create a new Domain; selecting another Domain will create a sub-Domain.



2. Enter the name of the Domain you wish to create and click **OK**.



3. Your new Domain will appear in the Nodes View.



Domain Information

The available information for a domain can be viewed using the following tabs:

- “Nodes”
- “Edit Parameters”
- “Storage”
- “Policy”
- “Monitoring”
- “Properties”

Nodes

The **Nodes** tab shows the [list of the nodes](#) in the domain. Using this tab, you can also stop or restart the nodes in the domain.

Edit Parameters

The **Edit Parameters** tab allows you to modify the description of the domain.



Note

You cannot modify the name of a domain once it has been created.

Storage

The **Storage** tab provides an overview of the storage resources associated with the current Domain.

At the top of the tab are fields listing the number of the various types of storage resources in the Domain. These are:

<i>Number of Pools</i>	<i>Disk Pools (Local/Shared)</i>
<i>Space Used</i>	<i>Pack Pools</i>
<i>Volume</i>	<i>Archive Pools</i>

Below that is a table providing details for each individual Pool.

Figure 6.1. Storage tab

Nodes \ Edit Parameters \ **Storage** \ Properties \

Pools7

Space Used18.14GB

Volume155.33GB

Local Disk Pools3

Shared Disk Pools1

Tape Pools (pack)0

Archive Pools3

Summary

Name ^	Cache	Tapes	Partitions	Size	Space Used	Data size	Auto e...	Node	Used Safes	State
Disk Pool 1	No		2	36GB	11.34% - 4.08GB			node-01	142	On-line
Disk Pool 2	No		2	36GB	9.68% - 3.48GB			node-02	120	On-line
Disk Pool 3	No		2	33.65GB	0% - 0Bytes			vcentos	0	On-line
Shared Test	No		2	30GB	11.24% - 3.37GB			node-02*, vc...	116	On-line
Archive Pool		3		14.76GB	45.03% - 6.65GB	6.77GB	Yes	node-01, no...	11	On-line
Export Pool		0		0Bytes	0% - 0Bytes	0Bytes	No		0	On-line
On-demand		1		4.92GB	11.25% - 567MB	970.05MB	No	node-01, no...	3	On-line

Policy

The **Policy** tab lists all the storage policies that have been configured for the disk pools contained in the domain.

Monitoring

You can monitor the activity and performance of the nodes in the domain using the **Monitoring** tab. For more information about node monitoring and a description of this tab, please see the topic [“Node monitoring”](#) in [Chapter 5, The Circle](#).

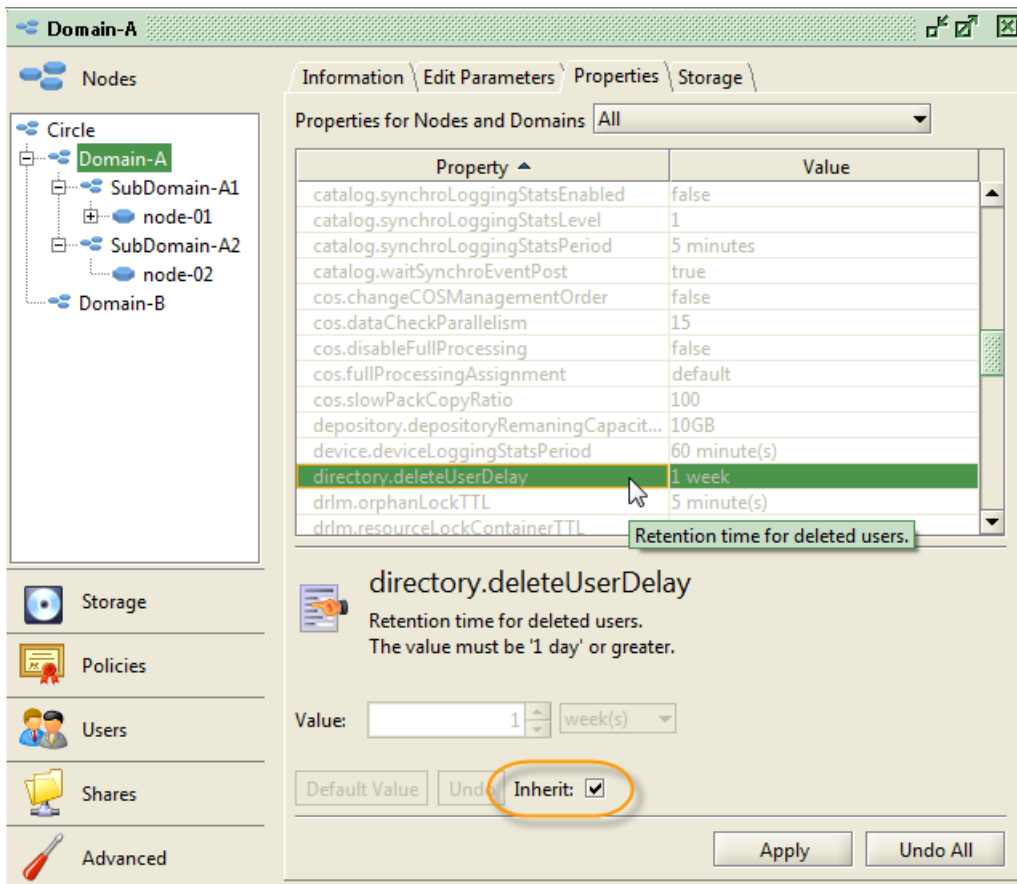
Properties

By default, the properties set at the level of the Circle are inherited by all the elements — the Domains, sub-Domains, and Nodes — within that Circle. It is, however, possible to override this and set properties specific to a Domain or to a Node.

The properties displayed in this tab are the same as the [“Circle Properties”](#). In most cases, the properties here are greyed-out because they are inheriting their values from the Circle, which is the default behavior.

You can uncheck the **Inherit** check box associated with some of the properties and then make a change to the value of the property. For some properties it is only possible to change the value for the entire Circle, in which case it will not be possible to uncheck the **Inherit** check box.

Figure 6.2. Properties tab



Tip

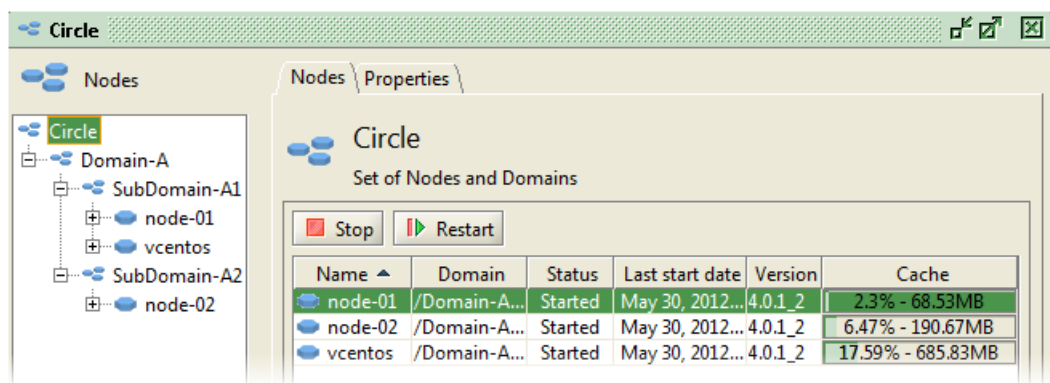
It is strongly advised that you consult with the Active Circle support team before making changes to any property values.

Chapter 7. Nodes

The *Node* forms the core of the Active Circle file storage system. In its simplest form, it is the representation in the system of your server and its attached storage. This chapter describes how to configure Active Circle nodes through the Administration Tool interface.

You can get a quick overview of the nodes in your configuration by going to the **Nodes** view and selecting the **Circle** item or any of the domains under it. In addition to the name and domain of each nodes, the table also shows information about the status, start date, software version and cache of each node.

Figure 7.1. Node List



In this panel you can also stop and restart the nodes. These operations are described further down in this chapter.

Adding a New Node

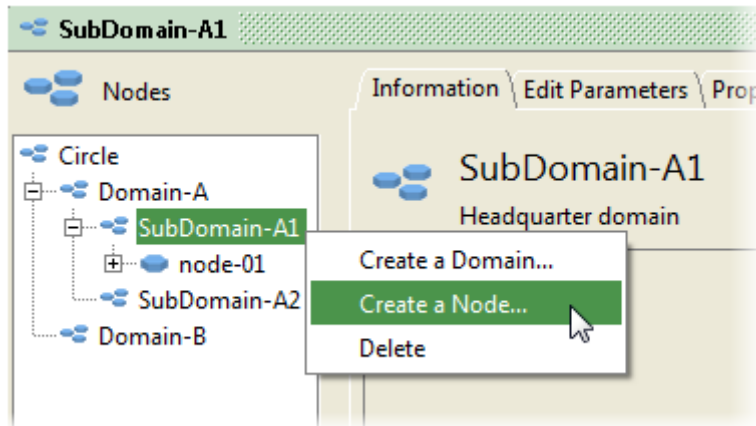
The first node in a Circle is created as part of the installation process. Any subsequent nodes must be created either in the Administration Tool or through the Active Circle Command Line.

Procedure 7.1. Creating a new node


The process for creating a new node is done in 3 steps:

1. Creation of the node in the Administration Tool
2. Installation of the Active Circle software on the server hosting the node. For instructions on how to install the software, please refer to the *Active Circle Installation Guide*.
3. Starting up the Active Circle service on that node.

To add a new node through the Administration Tool, right-click on the domain or subdomain where you want to locate it and click on **Create a Node...**




A dialog box opens where you can enter a **Name** for the node to be created.



By default, the option **Add the Node to the Policy Management** is checked. If unchecked, the node will be removed from the pool of machines used by the Active Circle service to distribute the work of internal system tasks. For more information, please see [Chapter 27, Policy Management](#).

Click the **OK** button to create the node.

A dialog box will appear displaying the identification key for the newly created node. You will need this key when installing the Active Circle software to the server for this node. To copy the key, right-click on the key in the dialog box and select **Copy**.

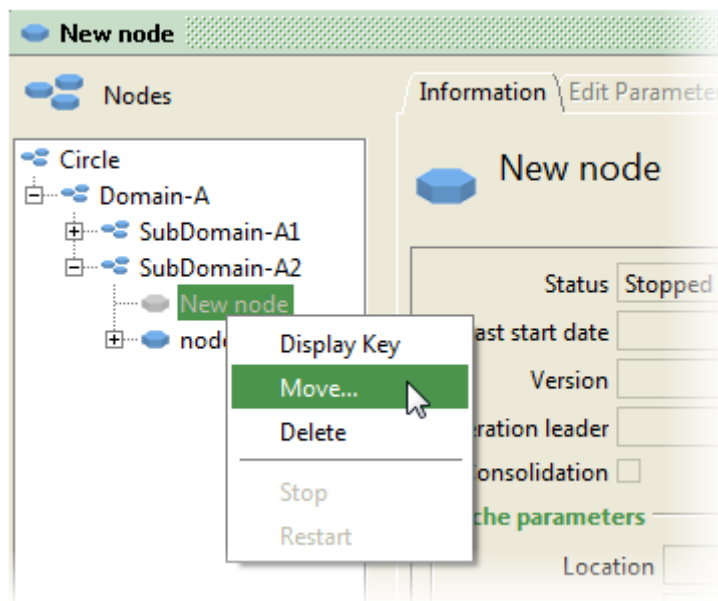


Tip

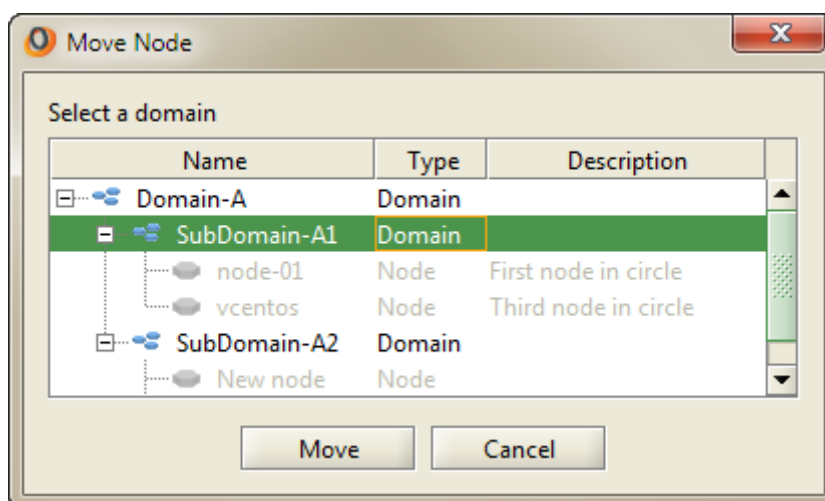
You can access the identification key after the installation by right-clicking on the greyed-out node in the **Nodes** view and selecting **Display Key**. The dialog box showing the key will display. You can then right-click on the key in the dialog box and select **Copy**.

Moving a Node

It is possible to move a newly created node to another domain using the Administration Tool. Right-click on the greyed-out node in the **Node** view and select **Move...** from the menu.



The **Move Node** dialog will be shown:



Select a domain in the list, and click on **Move** to place the node in another domain. The node will be moved in the tree view displayed by the **Node** view.



Note

The move operation is only possible on a node where the Active Circle software has not yet been installed. Once the Active Circle software has been installed on the server and the service started, the node can no longer be moved between domains.

Installing and Starting the Node

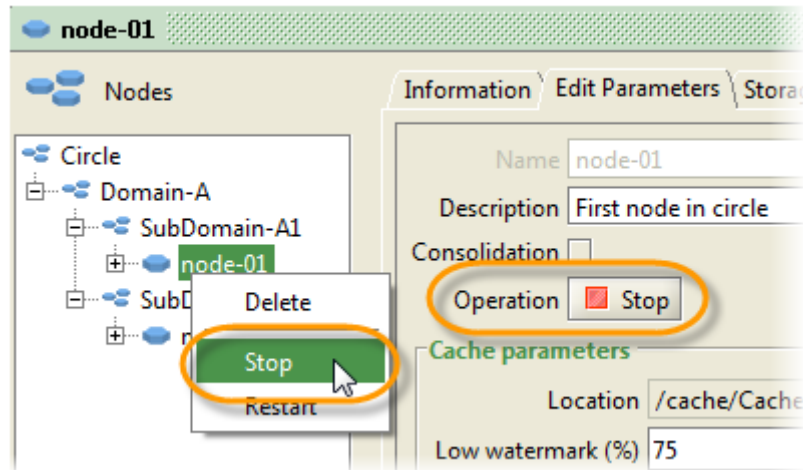
Your new node will appear in the **Nodes** view as greyed-out. It will remain in this state until you install the Active Circle software and [start up the Active Circle service](#) on that server. Please refer to the *Active Circle Installation Guide* for installation instructions.

Stopping a Node

You can stop the Active Circle service either through the Administration Tool or through the command line (see [“Stopping the Active Circle Service”](#)).

There are two ways to stop the service through the interface. After selecting the node in the **Nodes** view, you can either right-click on it and select **Stop**, or you can click on the **Edit Parameters** tab and then click the **Operation Stop** button.

Figure 7.2. Service Stop Menu Item and Button



You can also stop a node using the **Stop** button on the [Node List](#).



Important

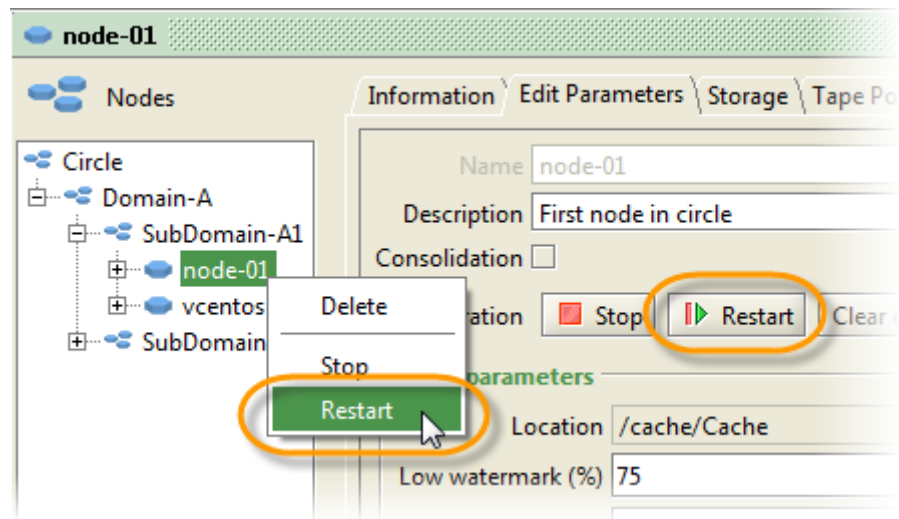
A node that is not running, cannot be started from the Administration Tool interface. This means that if you stop the node as described using the Administration Tool, you cannot start it again from here. For information on starting the Active Circle service, see [“Starting the Active Circle Service”](#).

Restarting a Node

You can restart an Active Circle service that is already running using the Administration Tool.

To restart the service, right-click on the node to be restarted in the **Nodes** and then click on **Restart** on the node context menu. Alternatively, you can select the node item and then click on the **Restart** button on the **Edit Parameters** tab. The Active Circle service will be restarted on the node.

Figure 7.3. Service Restart



You can also restart a node using the **Restart** button on the [Node List](#).



Note

The **Restart** menu item is only available when the Active Circle service is already running on the node. If the node is not started, the menuitem is greyed out.

Removing a Node

The process of removing a node from a Circle is done in two steps:

1. Stop the Active Circle service on the node.
2. Delete the node through the Administration Tool interface.

Deleting the Node

You must stop the Active Circle service before you can delete a node. To stop the node, follow the instructions [in the section above](#). In addition, for reasons of data integrity and security, a node cannot be deleted if:

- it is defined as a data location in a storage policy,
- it is used for policy processing of one or more shares,
- it is in the list of nodes used by the Lock Manager,
- it is the only node in a HA Cluster.

If you intend to proceed with the removal of the node, you must first resolve any such conflicts.

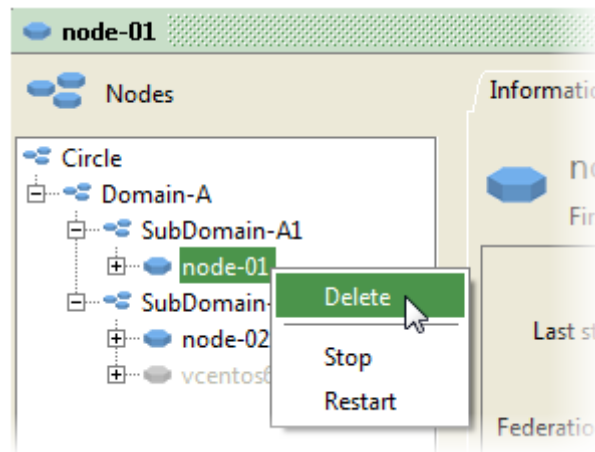


Warning

Deleting a node will delete all the Active Circle data in the storage pools that are attached to the node, even if the pools contain sole copies of the data. If you want to retain any of the data stored in the node's pools, make sure the data has been replicated to pools on other nodes first.

To delete the node, right-click on it in the **Nodes** view and click **Delete**.

Figure 7.4. Delete Node



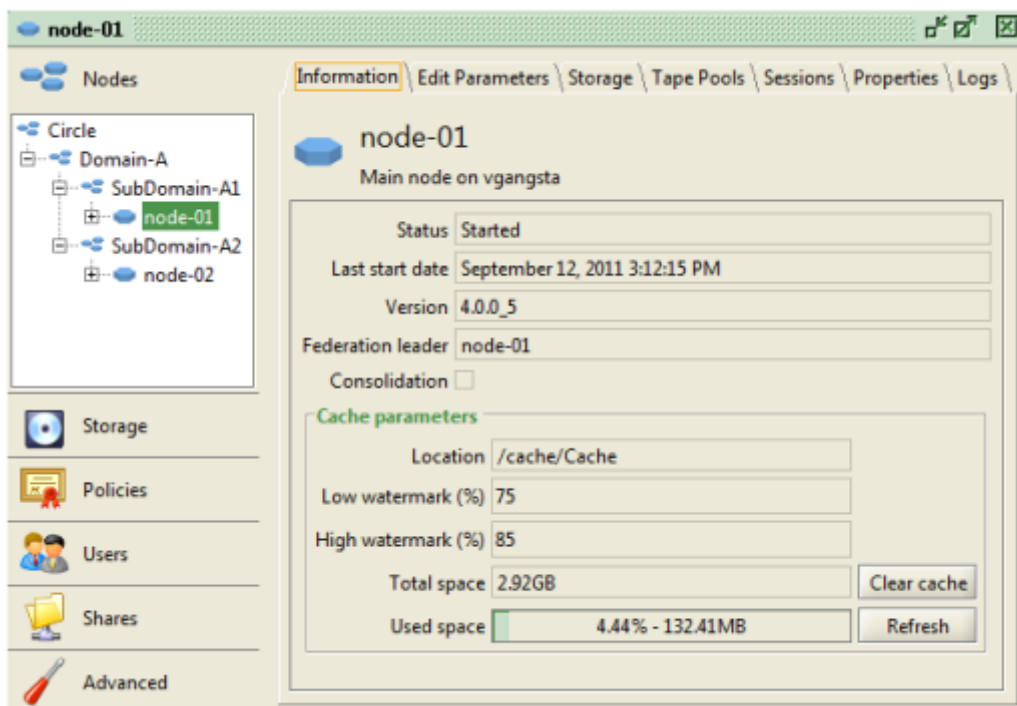
If you have met all the pre-conditions for deleting the node a message will be displayed asking you to confirm before deleting. Click **Yes** to delete, click **No** to cancel the operation.

If you have not met the preconditions, a message will be displayed informing you and you must correct the problem before trying to delete the node.

General Node Parameters

When you select a node in the **Nodes** view, a series of tabs are displayed. It is through these tabs that the nodes are managed. The first two tabs, **Information** and **Edit Parameters**, are where you can view and modify the general parameters for the node.

Figure 7.5. Node – Information tab



The node name is displayed at the top of the tab. The name is determined when it is created and may not be changed. The other fields are:

Description	A text field used to provide additional information about the node.
Status	The current status of the node.
Last Start Date	The date and time of the last start-up of the Active Circle service.
Version	The version of the Active Circle software running on the node. The last number (after the underscore) indicates the patch level.
Federation leader	The name of the node that has been elected Federation leader.
Consolidation	Activate consolidation for the current node. If checked, the node is a <i>Consolidation node</i> . Please see “Create a Consolidation Node” for more information.
Operation Stop button	This button will stop the Active Circle service on the server where the selected node is running. There is no warning message. If you use this button, you will have to restart the service manually on the server itself.



Warning

If you press the **Stop** button for the node to which you are connected, you will be disconnected from the Active Circle service. There is no warning message and you will have to access the server itself to restart the service manually. You cannot restart the service from a client using the Administration Tool.



Tip

Any changes made to the information in these fields must be confirmed by clicking the **Apply** button.

Cache

The node *Cache* is a temporary storage area used for improving file-access performance in the Active Circle system. When a user adds a file to a share, Active Circle first writes that file to the cache before processing it and transferring it to its proper location in the storage system as determined by the policies.

Similarly, if a user wants later to access that file, the system will first look in the cache for it. If it finds the file, it delivers it to the user from there. If the file is not in the cache, the system will then retrieve the file from its storage location, write a copy to the cache, and then serve up the file to the user from there.

Every node in a Circle has a cache directory that is assigned during the installation process. Generally, a separate, dedicated partition is used for this purpose.

Cache Parameters

General settings for the cache that cannot be modified:

Location	The path to the local cache directory on the server.
-----------------	--

Total space	Displays the total amount of local disk space available in the cache directory on the node.
Used space	Displays the amount of space used both as a percentage of the Total space available and as total volume used.

The two node cache parameters that can be modified are:

- **Low-water mark**
- **High-water mark**

These two settings allow you to control how much data is stored in the cache. The **High-water mark** represents the upper limit of storage for the node. When reached, the system will automatically search through the cache for the oldest files and delete enough of them to bring the node's storage utilization down to its **Low-water mark**.

The goal is to have the most frequently consulted files available in the cache, from where they can be served up to end users more rapidly, thereby improving system performance.

The **Refresh** button queries the local system and updates the display accordingly.

Clearing the Cache

The **Clear Cache** button allows you to purge the cache manually of its contents. You can find this button both on the **Information** tab and the **Edit Parameters** tab for the node. This operation will delete *all* of the files in the cache except those added within the current cycle.¹



Note

The cache purge will not delete files that do not exist in at least one storage pool (disk or tape).

Create a Consolidation Node

The Active Circle system stores the metadata for a node in the Catalog. Each node has its own catalog. To guarantee their integrity, these catalogs can be copied to one or more of the other nodes in the Circle. For example, in a Circle of five nodes, each node's catalog might be copied to three other nodes in the Circle. By default (and design), no single node contains copies of all of the other nodes' catalogs.

It is, however, possible to achieve this by designating one of the nodes as a *Consolidation* node, thereby allowing you to back-up all the catalogs for the Circle from a single machine.

To create a Consolidation node, go to the **Edit Parameters** tab and select the **Consolidation** check box and click **Apply**.



Important

While it is possible to have more than one Consolidation node (that is, the system will allow it), it is not necessarily advisable. Because of the disk activity and network traffic required by consolidation, having multiple Consolidation nodes can negatively impact performance.

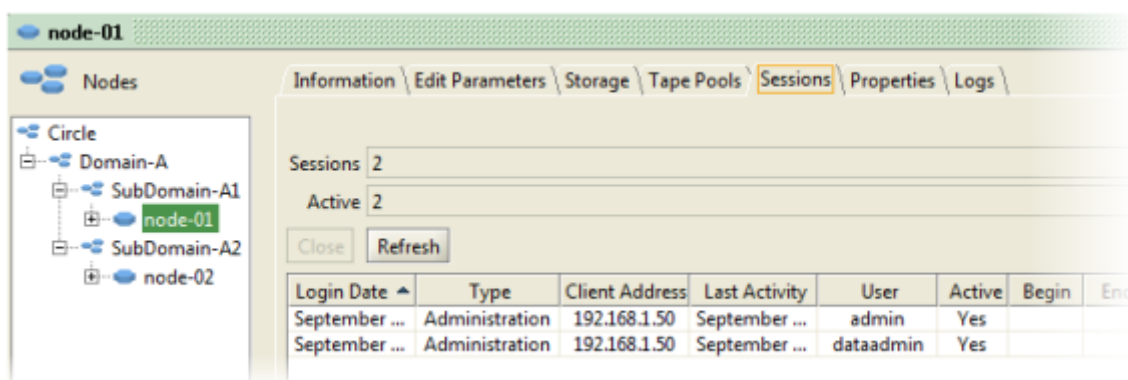
¹The contents of a cycle are all those files added within a given 30-minute period.

Sessions

The **Sessions** tab displays a table providing an overview of user activity in the system. The columns displayed by default are:

Login Date	User
Type	Active
Client Address	Begin
Last Activity	End

Figure 7.6. Sessions tab



The **Close** button allows an administrator to terminate a user's session.

The **Refresh** button refreshes the information displayed in this tab.

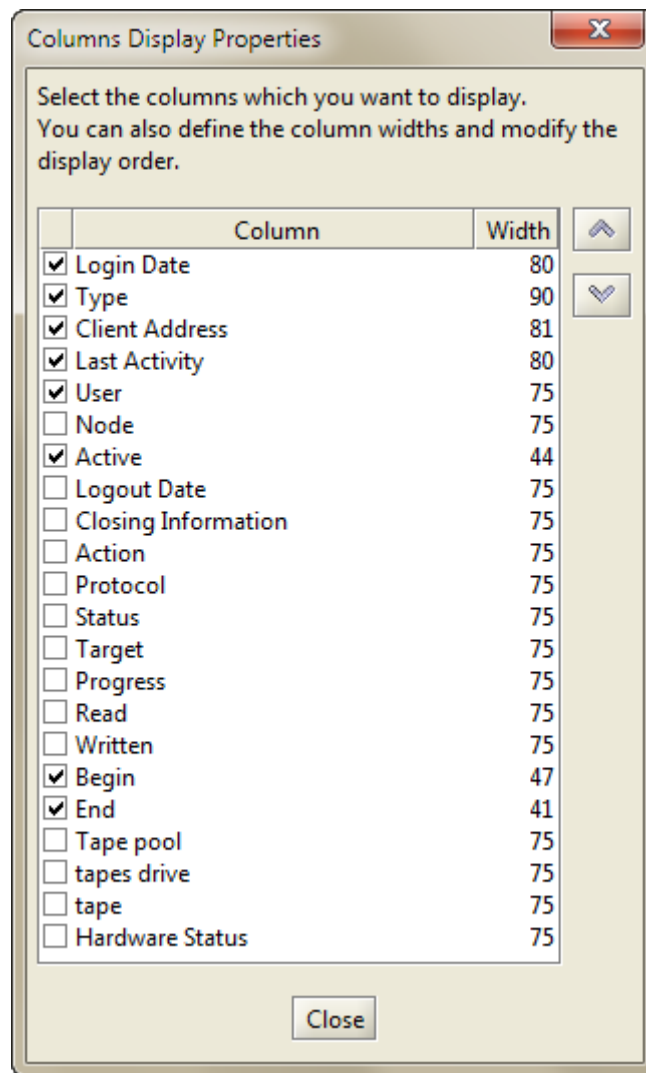
Modifying the Contents of the Sessions Table

You can change what information is displayed in the **Sessions** table in order to display the information that is most relevant to the particular type of activity that you wish to monitor. The properties that can be displayed are:

Name	Active	Read
Description	Logout Date	Written
Login Date	Closing Cause	Begin
Type	Action	End
Client Address	Protocol	Pool
Last Activity	Status	Drive
User	Target	Tape
Node	Progress	Hardware Status

To change the way columns are displayed, right-click on the column header row in the table and click **Column properties....** A dialog box pops up with a list of all the possible properties. To display a given property, check the box next to its name. Because these changes do not modify system data, they do not require validation: just close the dialog box.

Figure 7.7. Sessions table columns

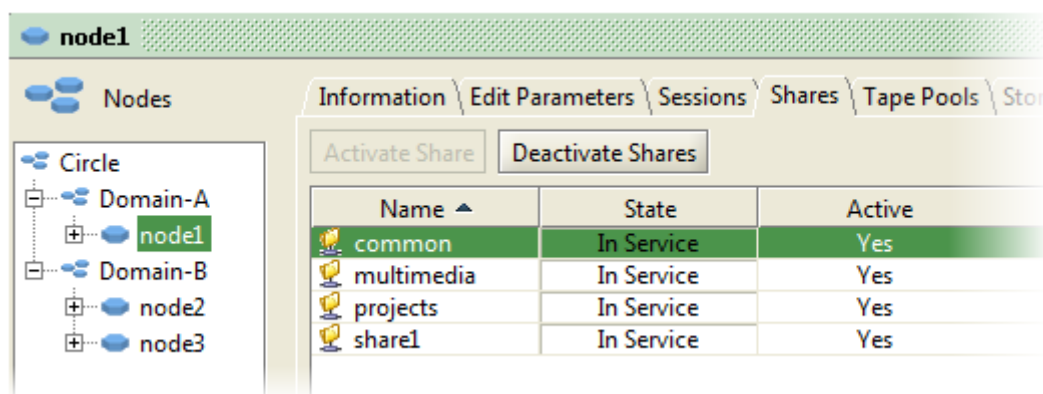


You can also filter the table contents by any character by using the [filter button](#) which appears in the upper right corner of the table if the session list is long enough to require vertical scroll bars.

Share association

The **Shares** tab lists all the shares in the circle and their status on this node. You can use the buttons above the list to activate or deactivate the share on the node. The **State** column shows if the share is started and in service. The **Active** column shows if the share is activated ("Yes") or deactivated ("No") on the node.

Figure 7.8. Shares tab



Share activation and deactivation can also be performed using the **Nodes** tab for a selected share in the **Shares** view. For more info, see [“Activation on Nodes”](#) in [Chapter 23, Shares](#).



Note

When a share is created, the system will automatically activate it on all the nodes of the circle.

Storage and Tape Pools

The **Storage** and **Tape Pools** tabs allow you to view and perform limited management tasks for the various storage devices attached to a node: disks and their partitions, libraries, tapes and tape pools. Since these tabs provide access to information and tasks from the perspective of the node, they can be useful for example when adding a new node to an existing configuration where shared devices and pools are already available.

Storage tab

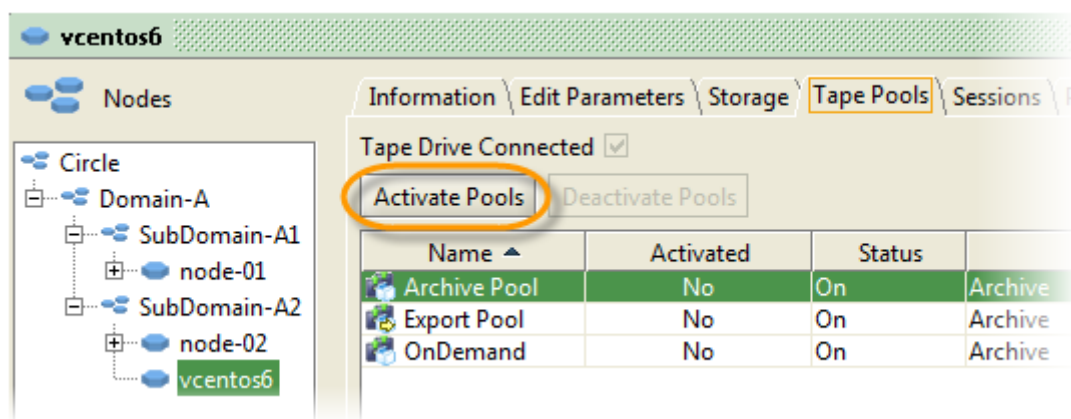
The **Storage** tab consists of three sub-tabs:

- Local Disks** This tab shows a list of the declared storage partitions on local disk storage attached directly to the selected node. You can use the buttons above the list to [Add Partitions](#), [Delete](#) partitions and [Create Pools](#) consisting of the partitions listed here.
- Libraries** This tab shows a list of tape libraries the node has access to, either local or shared.
- Tape Drives** This tab shows a list of the tape drives available to the node. Using the buttons above the list, you can [Clean](#) drives, [Delete](#) drives, and place drives under maintenance.

Tape Pools tab

The **Tape Pools** tab shows a list of tape pools available to the selected node. If the tape pools belong to a shared library, you can use the [Activate Pools](#) button to activate one or more of the pools on this node. This is convenient if you have just added the node to a configuration and you want it to have access to all or some of the shared tape pools of the circle.

Figure 7.9. Node – Tape Pools tab



Managing Storage Devices and Pools

The full functionality for managing storage devices and pools across the whole system can be found in the **Storage** view. When setting up storage devices and pools from scratch, it would generally be more practical to use the **Storage** view.



Note

The new location for managing all types of storage pools and devices in Active Circle is the **Storage** view. All the information and functionality on the **Storage** tab and **Tape Pools** tab for a node are duplicates of the same items in the **Storage** view, which is described in this manual starting with the chapter [Chapter 8, *Introducing Storage Pools and Devices*](#).

Both the **Storage** tab (and its subtabs) and the **Tape Pools** tab have equivalent tabs under the **Storage** view. The following table shows where in this manual each tab is described.

Table 7.1. Mapping of Nodes to Storage view tabs

Tab in Nodes view	Location in Storage view
Storage	Chapter 8, <i>Introducing Storage Pools and Devices</i>
Local Disks	Chapter 13, <i>Disk Devices</i> Chapter 9, <i>Local Disk Pools</i>
Libraries	Chapter 14, <i>Library Devices</i>
Tape Drives	"Tape Drive Operations"
Tape Pools	Chapter 11, <i>Tape Pools</i>

Policy

The **Policy** tab lists all the storage policies which have been associated with the disk pools of this node.

Network Interfaces

The **Network Interfaces** tab for a node shows information about all the active network interfaces on the server the node is running on. For each interface, the name, description and IP address (and type) is shown.

Figure 7.10. Network Interface Information for a Node

Information \ Edit Parameters \ Storage \ Tape Pools \ Sessions \ Properties \ Logs \ Network Interfaces		
Name ▲	Description	Address
eth0	eth0	192.168.1.73 Unicast (vgangsta)
lo	lo	127.0.0.1 Unicast (localhost.localdomain)

Properties

The behavior of a node can be controlled in two ways: by modifying one of the lines in the **Properties** tab (which can be done at the level of the Circle, the domain or the node), and/or by making a change to the file `/activecircle/.localvars`. This section describes modifying the **Properties**. For information about modifying the `.localvars` file, see [Appendix A, Local Variables](#).

To access the node properties, click to select a node in the **Nodes** view and then click on the **Properties** tab.

The properties for a node are the same as those of the Circle (see “[Circle Properties](#)”) and by default all the nodes in a Circle inherit their properties from their domain. It is possible, however, to modify the **Properties** for an individual node by selecting it in the **Nodes** view and clicking on the **Properties** tab.



Note

Not all properties can be modified at the node level and can only be modified at the level of the Circle. These properties will not have the **Inherit** check box.

Please see “[Editing Properties](#)” in [Chapter 5, The Circle](#) for more information about the properties and how to edit them.

Logs

Active Circle has an extensive collection of logging services that provide a detailed view of the internal activity of each node. These services can be configured through the Administration Tool interface in the node **Logs** tab.

Figure 7.11. Logs Tab

Information | Edit Parameters | Storage | Tape Pools | Sessions | Properties | **Logs**

Show external log services ☐

Name ▲	Activated	Level	Filter
ACL.controls	Yes	OFF	
ACL.netProtocols	Yes	OFF	
Admin	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Api	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Archive	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Box	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Catalog	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Catalog.service	Yes	WARNING	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Cluster	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Command	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Configuration	Yes	CONFIG	FATAL, ERROR, WARNING, INFO, CONFIG, ...
DataSrc	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Depository	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Destaging	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Device	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Federation	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Key	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Lease	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Medium	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Plug	Yes	CONFIG	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Remote	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
RLM	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
RLM.client	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
RLM.dispatch	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
RLM.lookup	Yes	WARNING	FATAL, ERROR, WARNING, INFO, CONFIG, ...
RLM.stats	Yes	WARNING	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Security	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...
Service	Yes	INFO	FATAL, ERROR, WARNING, INFO, CONFIG, ...

Log Files Management

Max log file size: 10MB Modify

Max log file count: 100 Modify

Log retention period: 3 months Modify

Viewing Node Logs

The log files for a node are stored in a dedicated directory on each node:

```
/activecircle/cell/data/Log/
```

The current log file, which is the file Active Circle is always writing data to, is named `lastcell.txt`. When the log file reaches its maximum allowed size, it is copied to an archive file and its old content is purged.

The archived log files are stored in a compressed format with a file name representing the date and time of the creation of the file: `20110105104835615.txt.gz`.

In the context of a troubleshooting operation, you may be asked to send the log files for a node to an Active Circle support representative. To do so, proceed as follows:

Pack the log files into a tar file using the command:

```
tar czf /tmp/logs_CellName.tar.gz /activecircle/cell/data/Log/
```

If you have an FTP account with Active Circle support, log into the [Active Circle FTP site](#) and copy the tar file to your project directory.

Node Log Administration

The **Logs** tab allows you to manage the various logging services running on the node, including external logs.

To display the external logs, tick the check box **Show external log services**.

The columns in the table are:

Log Service Name	Name of the logging service
Internal	This column only appears when the check box Show external log services is checked and indicates which log services are internal to Active Circle and which are external.
Active	Indicates whether or not the log service is running
Level	Indicates the level of detail of the log service and the type of information being collected.
Filter	Event filters applied to the log service

Level

It is possible to modify the level of detail for any individual log service. The default level for most logging services is `Information`. To troubleshoot a particular problem, however, you might switch the trace level to `All` to obtain a more detailed picture of what the system is doing.

To do this, double-click on one of the logs in the list. The **Log** dialog box is displayed. This dialog has two tabs, **General** and **Filters**.

Figure 7.12. Node Log Administration dialog box – General tab





Figure 7.13. Node Log Administration dialog box – Filters tab



The **Level** drop-down list allows you to choose from different categories and levels of logs. They are:

Off	Config
Fatal	Debug 1
Error	Debug 2
Warning	Debug 3
Info	All

The **All** option provides the highest level of detail.

Using the **Filter** options, you can uncheck any types of log you do *not* wish to include in the logs.



Note

For reasons of performance and practicality, it is not recommended leaving the log files at the **All** level. The recommended level for most logs during normal operation is **Information**.

Log file management

The **Logs** tab also provides functionality for managing the size, number, and life span of the log files.

Max log file size	Maximum size of a log file. When this size is reached, the log file is saved in compressed format and the contents of the <code>lastcell.txt</code> is deleted.
Max log file count	The maximum number of log files to store. Above this number, the older log files will be deleted as new ones are created.
Log retention period	Length of time to store a log file. When the age of log file exceeds this parameter, it will be automatically deleted.

These three parameters can be modified by clicking on the **Modify** button, entering a new parameter in the dialog box that pops up, and clicking **OK**.

Monitoring

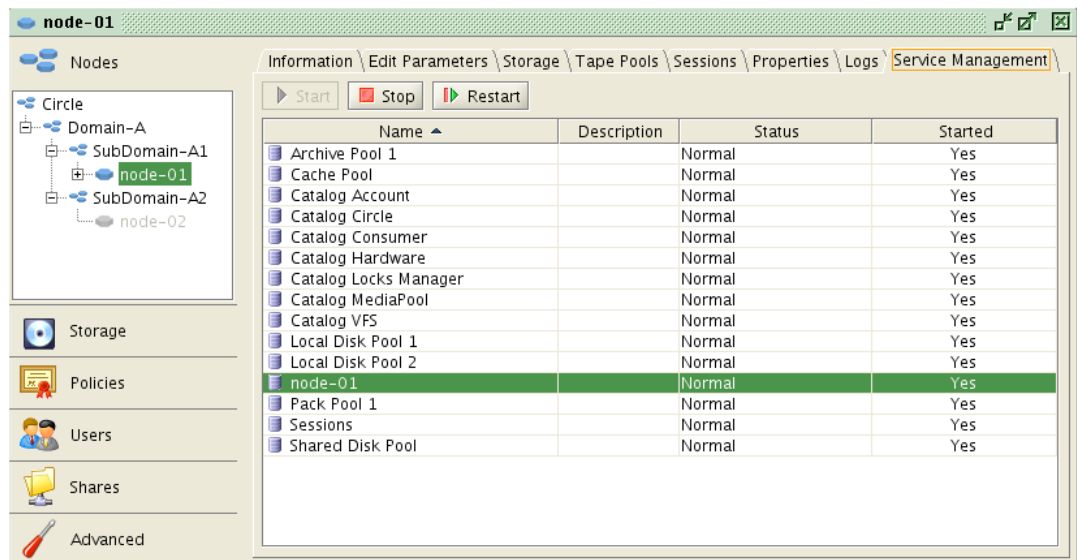
You can monitor the activity and performance of the node using the **Monitoring** tab. You can also monitor several nodes at the same time at the domain and circle level. For more information about node monitoring and a description of this tab, please see the topic “[Node monitoring](#)” in [Chapter 5, The Circle](#).

Service Management (Server Only)

The **Service Management** feature can only be accessed and run directly on a node server. The tab provides an overview of the services running on the current node in the form of a table displaying the following information:

Service Name	Status
Service Type	Started

Figure 7.14. Service Management tab



Three buttons - **Start**, **Stop**, and **Restart** - allow you to manage these services.

You can restart the Active Circle service on the current node by selecting the service `node_name` and clicking **Restart**.



Note

By default, the **Service Management** tab is hidden. It can be made visible by adding the property `-Dactivecircle.appli.ui.servicesTabVisibility=true` to the `.localvars` file. Note that this change will only be visible if you run the Administration Tool client directly on the server hosting the node and not by using the Java Webstart client.

For more information on setting local variables, please see [Appendix A, Local Variables](#).

Chapter 8. Introducing Storage Pools and Devices

Before a storage device can be used by Active Circle, it has to be included in a *Storage Pool*. In Active Circle, a *pool* is a grouping of storage *devices* that constitute a storage unit.

- A **Disk Pool** is a group of local disk partitions attached to the same node.
- A **Shared Disk Pool** is a group of shared disk partitions defined in the Circle.
- A **Tape Pool** is a group of tapes attached to libraries in the Circle. A tape pool will be shared automatically if it is activated on two or more nodes that have access to a shared tape library.
- A **Cloud Pool** is a storage and archiving space hosted in the cloud, off site.

The data in pools are managed through *storage policies* (disk pools and tape pack pools) or *archive policies* (archive tape pools and cloud pools).

The **Storage** view groups all the storage-related items in the Administration Tool interface. You can manage all your disk pools, partitions, tape pools, libraries and cloud pools from this single view. Some of the items and tasks in the Storage view are still accessible from the Nodes view, but the functionality is the same.

Figure 8.1. Storage View (the sub-categories are highlighted in blue)

The screenshot displays the 'Storage' view in the Administration Tool. The left-hand tree view shows the hierarchy: **Pools** (containing Disk Pools, Tapes, and Cloud) and **Devices** (containing Disks and Libraries). The right-hand pane is titled 'Pools' and shows summary statistics: 7 Pools, 7.68GB Space Used, and 4.73TB Volume. Below this is a 'Summary' table listing individual pools with their properties.

Name	Cache	Tapes	Partitions	Size	Space Used
Archive Pool		1		799.97GB	0% - 0Bytes
Cache Pool	Yes		1	5GB	51.11% - 2.56G
Disk Pool 1	No		2	13GB	0.13% - 16.99M
Disk Pool 2	No		1	5GB	51.11% - 2.56G
Export Pool		2		1.56TB	0% - 0Bytes
Pack Pool		3		2.34TB	0% - 0Bytes
Shared Pool	No		1	18GB	14.2% - 2.56G

The **Storage** view is divided into the main categories **Pools** and **Devices**. Each category is again divided into items based on the storage types of either **Disks** or **Tapes/Libraries**.

Storage Pools

The **Pools** item has tabs that show high level information for all the **Disk Pools**, **Tape Pools** and **Cloud Pools** in the Circle. You cannot change any information at this level of the view. The **Information** tab shows total numbers for all pools, and a list of all the disk, tape and cloud pools that have been configured.

Figure 8.2. Information tab

Name	Cache	Tapes	Partitions	Size	Space Used	Data size	Auto expand	Node	Used Safes	State	Availability	Status
Archive Pool		1		799.97GB	0% - 0Bytes	0Bytes	Yes	node-01, node-02	0	On-line	Normal	Normal
Cache Pool	Yes		1	5GB	51.11% - 2.56GB			node-02	84	On-line	Normal	Normal
Disk Pool 1	No		2	13GB	0.13% - 16.99MB			node-01	1	On-line	Normal	Normal
Disk Pool 2	No		1	5GB	51.11% - 2.56GB			node-02	84	On-line	Normal	Normal
Export Pool		2		1.56TB	0% - 0Bytes	0Bytes	Yes	node-01, node-02	0	On-line	Normal	Normal
Pack Pool		3		2.34TB	0% - 0Bytes	0Bytes	Yes	node-01, node-02	0	On-line	Normal	Normal
Shared Pool	No		1	18GB	14.2% - 2.56GB			node-02*, node-01	84	On-line	Normal	Normal

For each pool, depending on the type, the data columns display information indicating whether the pool is a cache pool, how many partitions the pool contains, how many tapes have been assigned to the pool, the size, whether auto expand has been activated, the number of used safes, node information and more. For cloud pools, no cache or partition information is applicable. For shared pools, the **Node** column lists all the access nodes. For shared disk pools, the leader node is indicated by an asterisk (*).

Tapes

If the Circle has been configured with one or more tape libraries, the **Tapes** tab will display a list of all the tapes attached to the system, whether they are part of configured Active Circle tape pools or the default pool of tapes available in the libraries. Total numbers for the tapes in the system are displayed above the list of tapes.

Figure 8.3. Tapes tab

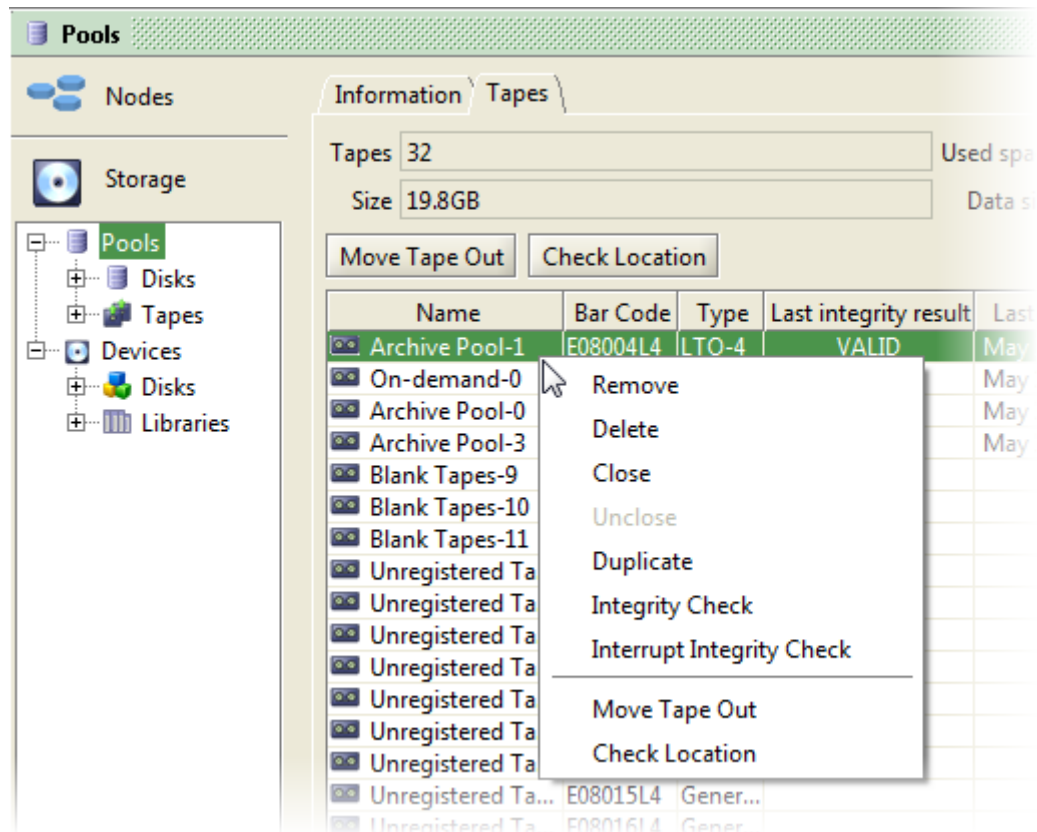
Name	Bar Code	Type	Size	Used Space	Data Size	Closed	Full	Location
Archive Pool 1-0	E01006L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 0 - Slot 1029
Archive Pool 1-1	E03006L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 1 - Slot 1029
Archive Pool 1-2	E01007L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 0 - Slot 1030
Archive Pool 1-3	E03007L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 1 - Slot 1030
Blank Tapes-0	E01001L4	LTO-2						Library 0 - Slot 1024
Pack Pool 1-1	E03001L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 1 - Slot 1024
Pack Pool 1-2	E01002L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 0 - Slot 1025
Pack Pool 1-3	E03002L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 1 - Slot 1025
Pack Pool 1-4	E01003L4	LTO-2	200GB	0% - 0Bytes	0Bytes	No	No	Library 0 - Slot 1026
Unregistered Tapes-5	E03003L4	Generic LTO						Library 1 - Slot 1026
Unregistered Tapes-6	E01004L4	Generic LTO						Library 0 - Slot 1027
Unregistered Tapes-7	E03004L4	Generic LTO						Library 1 - Slot 1027

You can perform several tasks on the tapes from this view. The buttons indicate actions that are available for all the tapes. **Moves Tapes Out** will move a selected tape from the currently

assigned slot to the library mailbox. Clicking the **Check Location** will update the location information for the tapes in the library.

To access more tape tasks, select one or more tapes in the list and right-click. A context menu will be displayed with the valid actions for the tape(s) you selected. If the tapes you selected belong to different tape pools, only the basic tasks indicated by the buttons above the list will be available.

Figure 8.4. Tape actions in pools



The same information is available when clicking directly on the **Disks**, **Tapes** or **Cloud** item in the tree under **Pools**. The list of pools is also displayed under the corresponding items in the tree view. If a pool element (local, shared, cloud or tape) in the tree view is greyed out, it means that the pool is not activated or any nodes or that none of the nodes it is activated on are started.

For information about how to create and manage Storage Pools, see the following chapters:

- [Chapter 9, Local Disk Pools](#)
- [Chapter 10, Shared Disk Pools](#)
- [Chapter 11, Tape Pools](#)
- [Chapter 12, Cloud Pools](#)

Pool Type Indicators



In the Administration Tool interface, the pool types are indicated by different graphical representations. Each type has its own icon, as shown in the table below. For each icon, one or more indicators can be added to communicate the state of the element the icon represents. Status indicators are explained within the appropriate sections for each element in this guide.

Table 8.1. Pool types

Icon	Description
	Local Disk Pool.
	Shared Disk Pool.
	Tape Pool (pack pool or default pool of tapes in the libraries).
	Archive Tape Pool.
	Export Archive Tape Pool.
	Cloud Pool.

There are several indicators that may be dynamically added to the icons in the interface. These work as modifiers that change the icon to present additional information (normal state, warnings, errors, informational) to the administrator. Two of these modifiers of the informational type have the same meaning for all pools, and are listed in the table below. For more information on other specific status indicators, see the section in this document describing the particular element it relates to.

Table 8.2. Global modifiers

Indicator	Description
	The pool is <i>Off-line</i> . This State can be changed on the Edit Parameters tab for the pool.
	The pool is <i>Read only</i> . This setting can be controlled using the Availability field on the Edit Parameters tab for the pool. This indicator can also mean that the pool is Full . In either case, no data can be written to the pool while it is in this mode.

Storage Devices

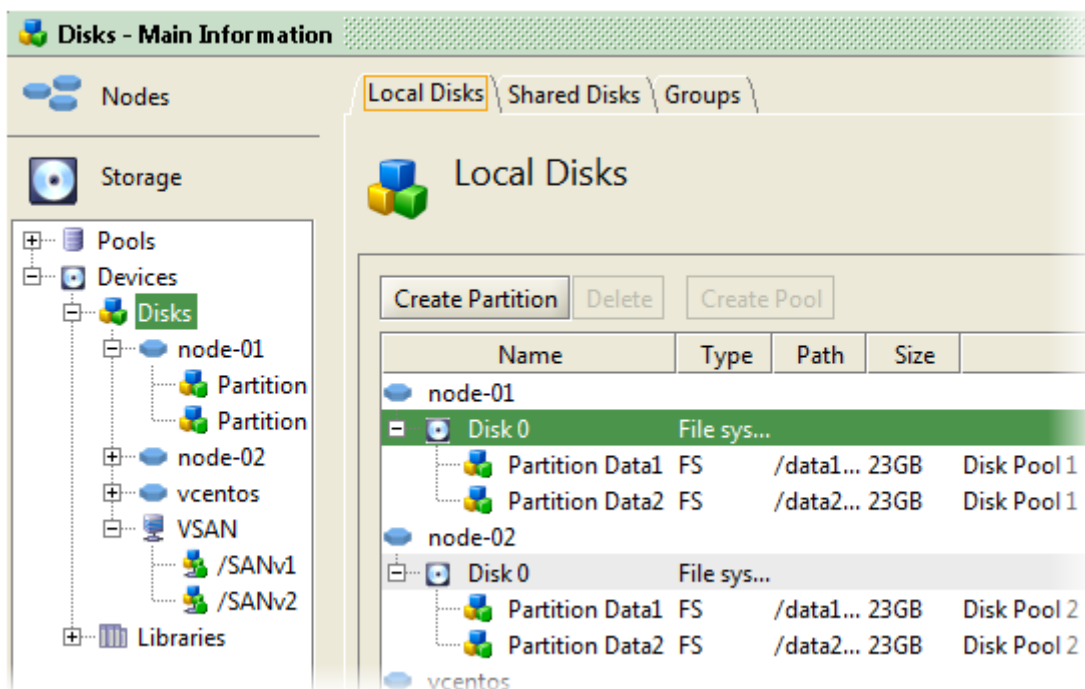
The **Devices** category contains information about the physical storage devices connected to the Circle. These are the devices that can be used by Active Circle storage pools.

Disks

The **Disks** item is used to create and manage both local and shared partitions. All the disks and partitions defined in Active Circle is attached to this item in the tree view. Click on the plus sign next to **Disks** to expand the view to show all the disk devices. This item has three tabs:

- **Local Disks** – Displays a list of all the disks and partitions directly attached to Active Circle nodes, and lets you create or delete partitions or create local disk pools.
- **Shared Disks** – Displays a list of all the shared partitions defined in SAN configurations, that one or more Active Circle nodes have access to. From here, you can also create or delete partitions and move partitions to a group.
- **Groups** – Lists the shared partition groups, and lets you create or delete groups.

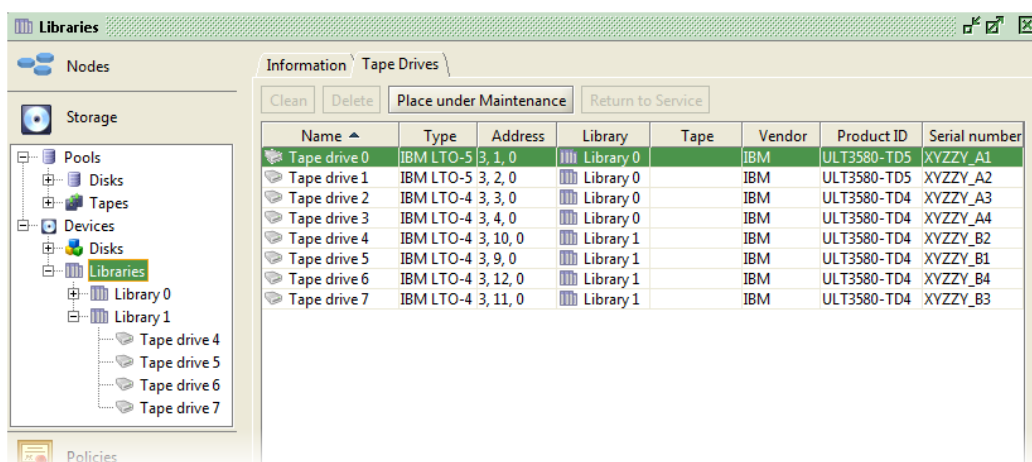
Figure 8.5. Disks – Local Disks tab



Tape Libraries

The **Libraries** item lists all the libraries attached to the Circle and their drives and tapes. From the **Tape Drives** tab you can perform library-specific tasks by using the buttons **Clean**, **Delete**, **Place under Maintenance**, and **Return to Service**.

Figure 8.6. Libraries – Tape Drives tab



For information about how manage Storage Devices, including creating disk partitions and manipulating tape drives, see the following chapters:

- [Chapter 13, Disk Devices](#)
- [Chapter 14, Library Devices](#)

Device Type Icons

In the Administration Tool interface, the device types are indicated by different graphical representations. Each type has its own icon, as shown in the table below. For each icon, one or more indicators can be added to communicate the state of the element the icon represents. Status indicators are explained within the appropriate sections for each element in this guide.

Table 8.3. Device Icons

Icon	Description
	Disk Partition
	Shared Partition (LUN)
	Shared Partition Group
	Disk
	Tape Library
	Shared Tape Library
	Tape Drive
	Shared Tape Drive
	Tape
	Node

Chapter 9. Local Disk Pools

A *Local Disk Pool* (sometimes referred to as a *Disk Pack Pool*) is a logical storage space associated with a node consisting of one or more *partitions* in which the Active Circle system stores file data. Before Active Circle can use a storage device, it has to be included in a pool.

In a typical configuration, one disk pool is created on each node, and all the partitions on the node's disk storage are included in the pool. The administrator can also choose to create several local disk pools on a single node, and assign each partition to a different pool. Please consult with Active Circle Support if your needs require this type of configuration. Cases where multiple disk pools should be created on a node include the following:

- A hierarchical storage configuration with different storage device types (for example using both SAS and SATA disk bays). In such cases one pool would be created for each storage type.
- If there is a need for storing several copies of the same data on a single node.



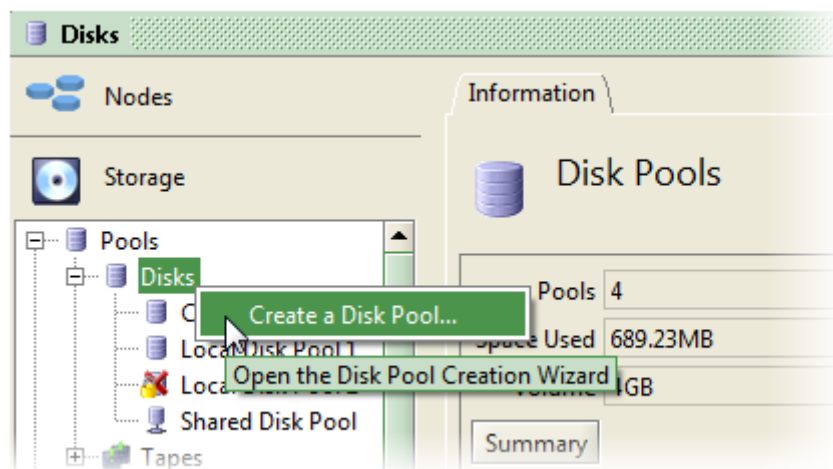
Note

A disk pool needs to be associated with a partition on a disk storage device. The partition is selected during the creation of the pool. If you have not created any partitions yet, you should do so before proceeding with creating the pool. For more information, see ["Adding a Local Partition"](#).

Creating a Local Disk Pool

To create a new pool, right-click on either the **Pools** or the **Disks** element at the top of the **Storage** view. Select **Create a Disk Pool** from the context menu. Alternatively, you can use the **Create** button on the **Information** tab for the **Disks** element. This will start the pool creation wizard.

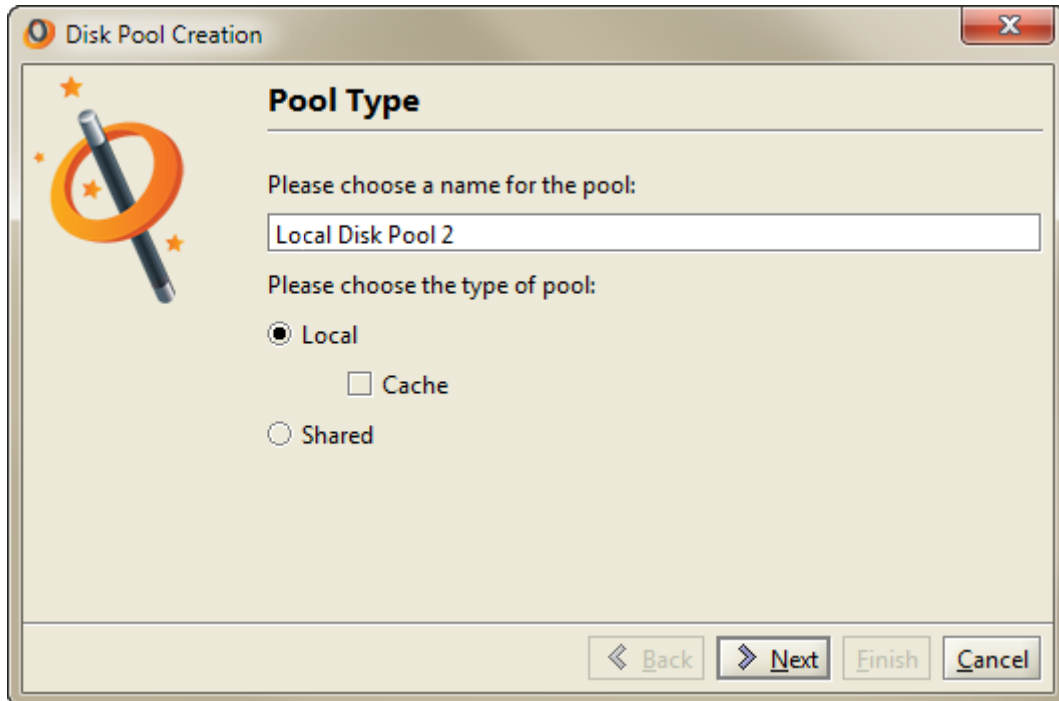
Figure 9.1. Starting the Pool Creation Wizard



In the first panel, you must enter a name for the pool and choose the type of pool to create. Select **Local**.

There are two other choices in this window. Place a check mark in the **Cache** box if you are creating a **Cache Pool**. This panel is also the starting point for creating a shared disk pool. For more information, see [“Creating a Shared Disk Pool”](#). Click on **Next** to continue.

Figure 9.2. Pool Name and Type



The screenshot shows the 'Disk Pool Creation' window with the 'Pool Type' tab selected. On the left is a graphic of a pen writing on a circle with stars. The main area contains the following text and controls:

Pool Type

Please choose a name for the pool:

Local Disk Pool 2

Please choose the type of pool:

☒ Local

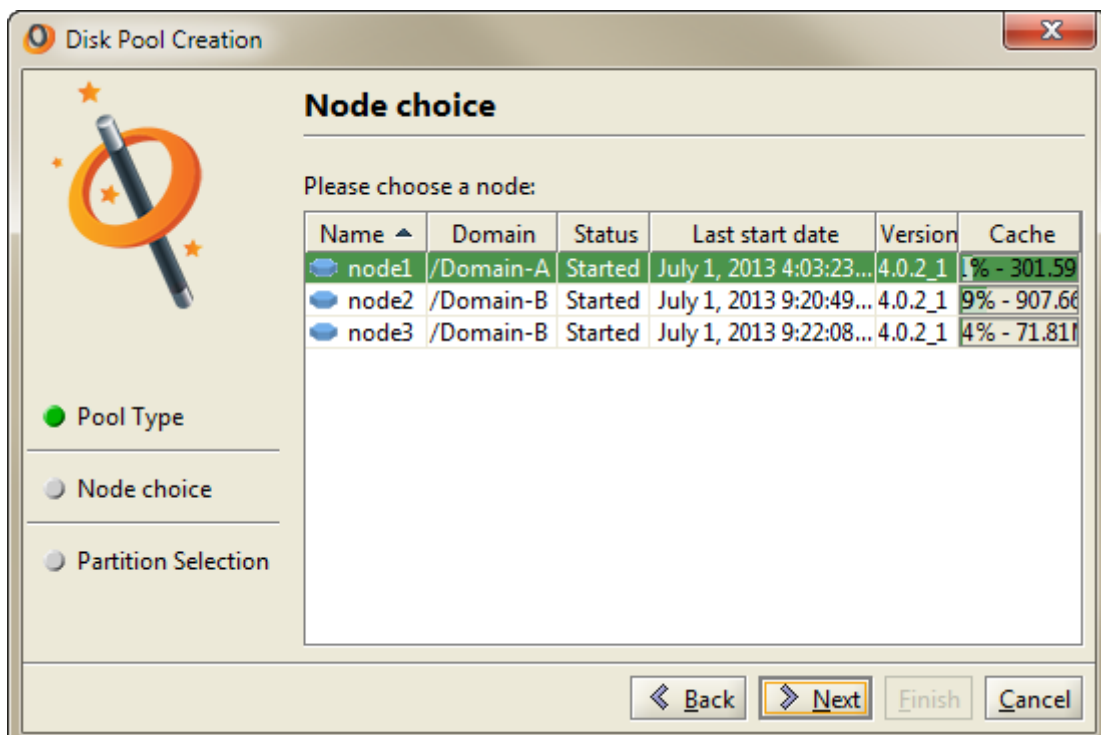
☐ Cache

☐ Shared

At the bottom are four buttons: Back, Next, Finish, and Cancel. The 'Next' button is highlighted.

Select the node where the disk storage you want to use for this pool is located. The list displays the available nodes in the Circle. Click on **Next** to continue.

Figure 9.3. Selecting the Node



The screenshot shows the 'Disk Pool Creation' window with the 'Node choice' tab selected. On the left is a graphic of a pen writing on a circle with stars. The main area contains the following text and controls:

Node choice

Please choose a node:

Name	Domain	Status	Last start date	Version	Cache
node1	/Domain-A	Started	July 1, 2013 4:03:23...	4.0.2_1	1% - 301.59
node2	/Domain-B	Started	July 1, 2013 9:20:49...	4.0.2_1	9% - 907.66
node3	/Domain-B	Started	July 1, 2013 9:22:08...	4.0.2_1	4% - 71.81

On the left side of the window, there are three radio buttons for navigation:

☒ Pool Type

☐ Node choice

☐ Partition Selection

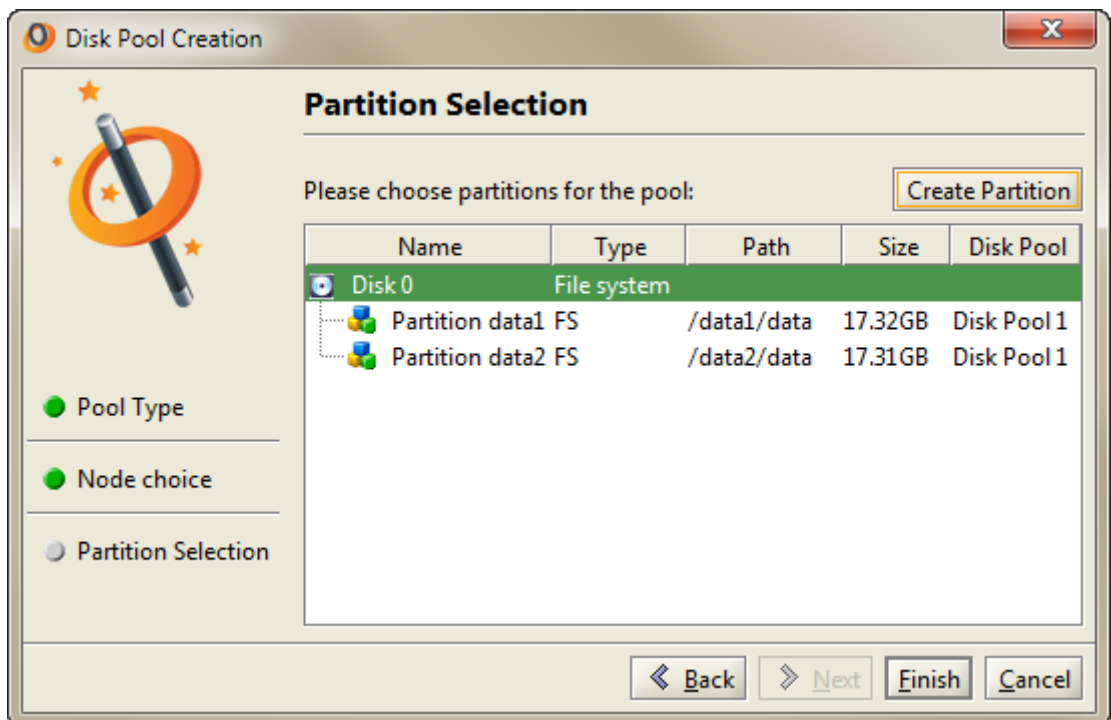
At the bottom are four buttons: Back, Next, Finish, and Cancel. The 'Next' button is highlighted.

The **Partition Selection** panel lists the disks and partitions on the storage attached to the node you selected in the previous panel. You must select one or more available partitions from the list (use the **Ctrl** key to make multiple selections) not already allocated to another pool, as indicated by the column **Disk Pool**. In a regular configuration where only storage device(s) of the same type are attached to the server, one would add all the partitions to the local disk pool.

If you want to add a partition that has not yet been declared in Active Circle, you can do it from this dialog by clicking on the **Create Partition** button. This will open the regular partition creation dialog, described in [“Adding a Local Partition”](#).

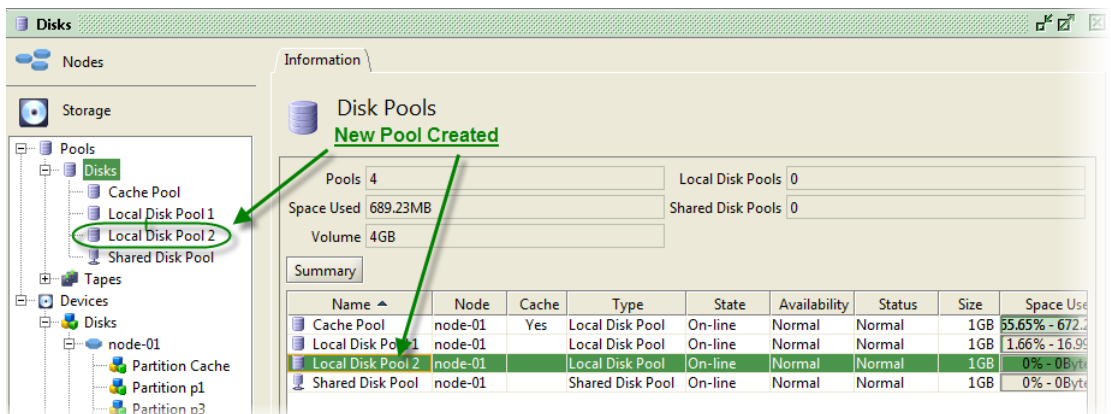
When done, click on **Finish** to create your pool.

Figure 9.4. Selecting the Partition



The newly created pool will be displayed as an element under **Disks** in the tree view, as well as in the list of pools under the **Information** tab for both the **Disks** and the top-level **Pools** element of the **Storage** view.

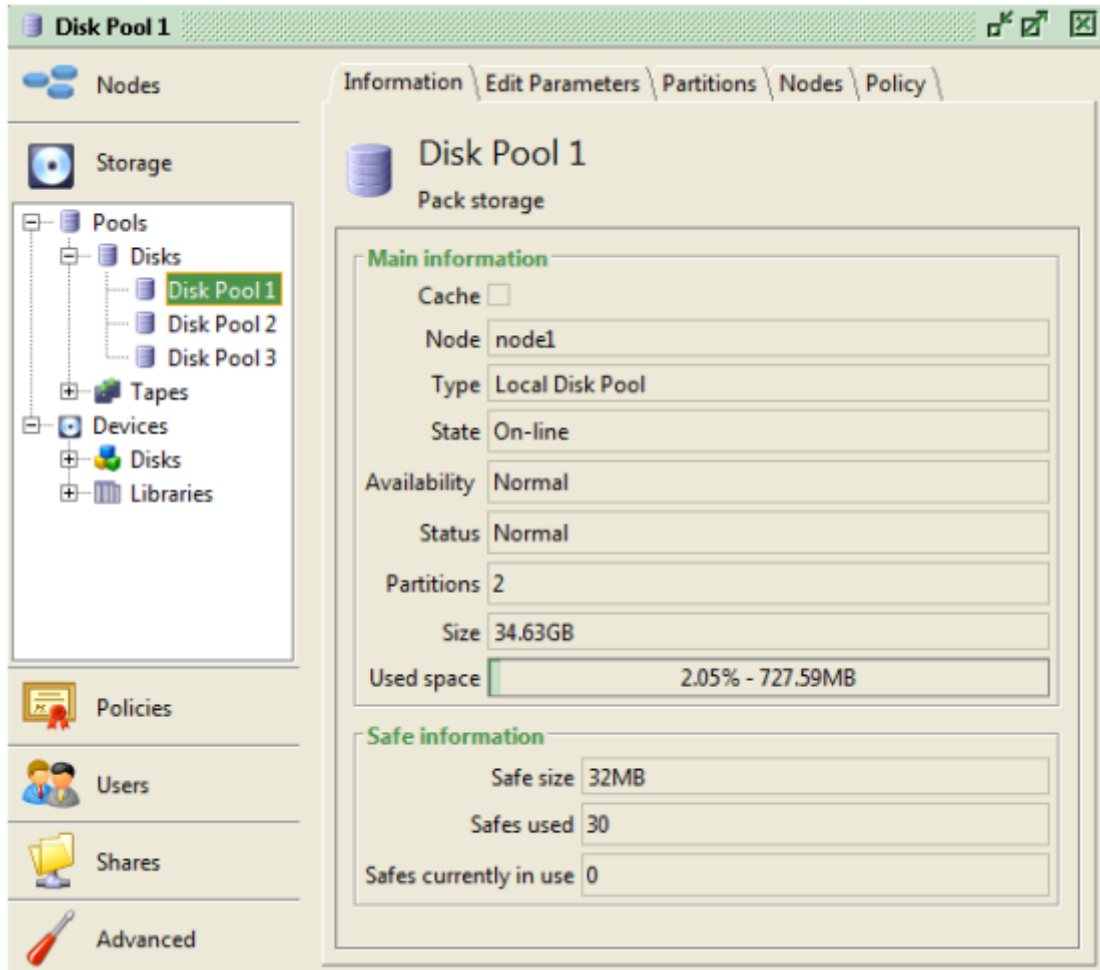
Figure 9.5. Pool Created



Pool Information

The **Information** tab displays all the information about the selected pool.

Figure 9.6. Pool Information tab



Main information

Cache	Check box to indicate whether this disk pool will be used as a cache.
Node	The node associated with this disk pool. More details on the node is available on the Nodes tab for the disk pool.
Type	Type (Local or Shared Disk Pool)
State	Current state of the disk pool. Possible values are: On-line, Off-line.
Availability	Current availability of the disk pool. Possible values are: Normal, Read only, Full, Error.
Status	Current status of the disk pool.
Partitions	Number of partitions associated with the disk pool.
Size	Size in GB of the disk pool.
Used space	Amount of disk space currently being used by the disk pool.

Safe information

Safe size	Size in MB of the Active Circle safes (also called packs). Defined during the installation procedure.
-----------	---

Safes used

Number of safes to which data has been written.

Safes currently in use

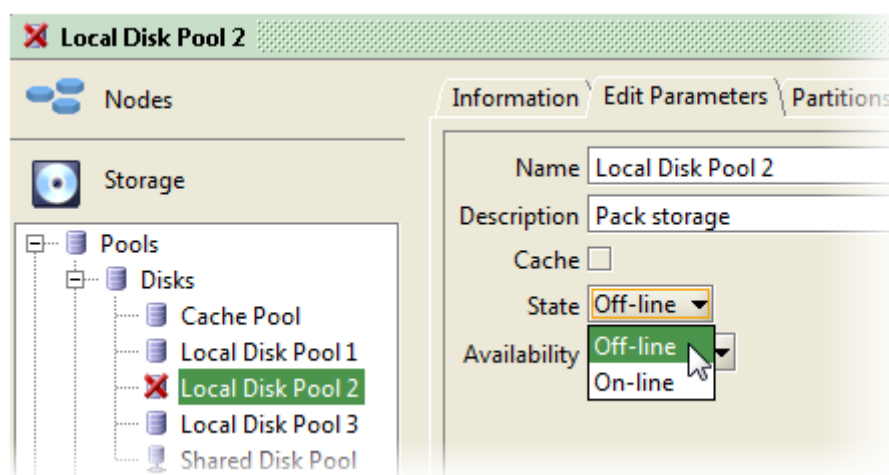
Number of safes currently open for use.

Disk Pool States

It is possible to modify the state of a pool by taking it off-line in order to prevent any read or write activity on the pool, something you might do during a maintenance operation on the server.

To take a pool off-line, in the **Edit Parameters** tab for the pool, click on the **State** drop-down list, select **Off-line**, and click **Apply**. A red 'X' appears on the pool in the tree view to indicate that it is off-line.

Figure 9.7. Off-line Pool



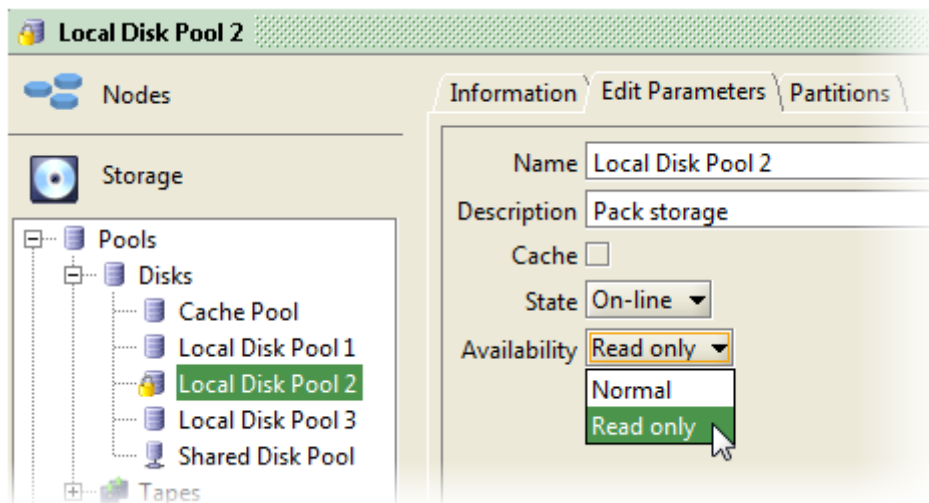
To return the pool to the on-line state, perform the same process selecting **On-line** and clicking **Apply**.

Disk Pool Availability

You can modify the availability of a pool by making it read-only.

To make a pool read-only, in the **Edit Parameters** tab for the pool, click on the **Availability** drop-down list, select **Read only**, and click **Apply**. This change is reflected in the **Availability** field that appears in the **Information** tab and in the yellow lock that appears next to the pool in the tree view.







Figure 9.8. Read-only Pool



Pool Status Indicators

Visual status indicators help identify the location of any potential issues affecting the local disk pools. The indicators are shown as small icons in the views pane, and as larger icons on the **Information** tab for the element. To find more information when a problem is indicated, please consult the supervision notes (see [Chapter 32, Supervision](#)).

Table 9.1. Local Disk Pool Status Indicators

Indicator	Meaning
	Local Disk Pool in normal operation, no problems detected.
	The pool is <i>Off-line</i> .
	The pool is <i>Read-Only</i> or <i>Full</i> . No more data can be written to the pool.
	There is an error with the pool.
	There is a problem with at least one partition in the pool. A <i>Supervision Note</i> has been issued.
	The indicators can be combined to show multiple problems. In this case, the pool has an error, it is off-line and one or more supervision notes have been issued due to a problem with the pool's partitions.

Cache Pools

In the **Create Disk Pool** wizard (see [Figure 9.2, "Pool Name and Type"](#)) there is a check box for creating a **Cache Pool**.

Cache Pools are used in configurations involving tape libraries. When a user makes a request for a file located in a tape library, Active Circle must first copy the *disk packs* (the Active Circle data storage format) that contain that file from the tape to a disk pool. If a cache pool exists, then it will be used. The system then extracts the file from there and copies it to the NAS cache, from where it will be delivered to the user who requested it.

Because of the additional step involved in working with files on tape, we recommend that you implement a cache pool on the node connected to the tape library.

If you want to convert a regular disk pool to a cache pool, you can go to the **Edit Parameters** tab of the pool and select the **Cache** check box. Then click **Apply** to activate the selection.

Figure 9.9. Selecting an Existing Pool as Cache Pool



Expanding Pools with New Partitions

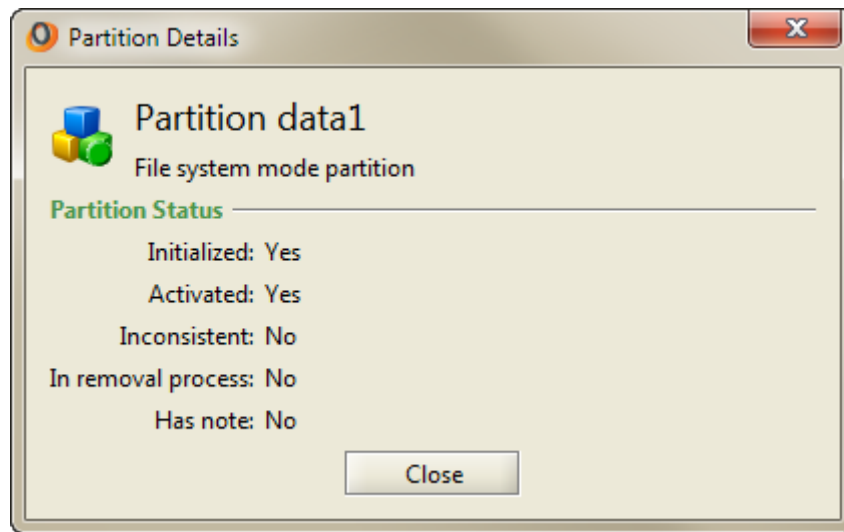
It is possible to augment the capacity of a node by adding a new partition to a pool without stopping the node and thereby interrupting the service.

You must first create the partition(s) you wish to add using the method described in [“Adding a Local Partition”](#). Once this is done, select the pool you wish to extend and click on the **Partitions** tab. Click the **Add Partitions** button, select the partition to add and click **OK**.

Partition Information

The **Partitions** tab of a disk pool shows a list of all the partitions in the selected pool. To view more information about a specific partition, select it in the list and click on the **Details** button. A dialog box will be displayed, showing status information for the partition:

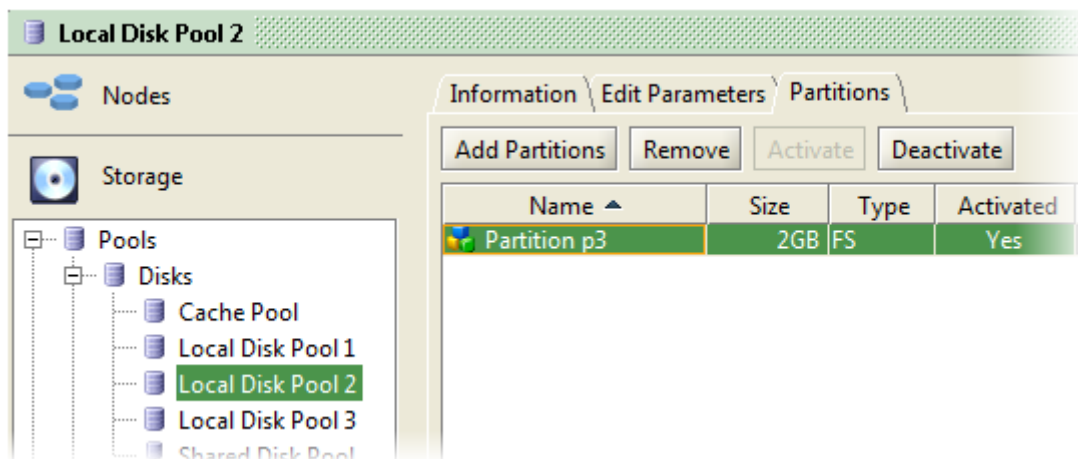
Figure 9.10. Partition Details



Removing or Deactivating a Partition

You can manage partitions from the **Storage** view. Select the pool containing the partitions and click on the **Partitions** tab. Select a partition and use the buttons to **Deactivate** or **Activate** the partition. When a local partition is deactivated, no data will be written to it and the data on it will not be available. The icon for the partition is greyed out while it is deactivated. If you reactivate the partition by clicking on the **Activate** button, the partition will again be available for reading and writing data.

Figure 9.11. Disk Pool — Partitions tab



If you deactivate or remove all the partitions of a pool, the pool itself will no longer be available. The availability status will be set to *Full* (a yellow padlock will be shown on the disk pool icon).

Removing a local partition will delete its contents. To remove a partition from a pool, first verify that the pool is on-line, as indicated by the **State** field on the **Information** tab for the pool. Partition removal will not be processed correctly if the pool is off-line. Then select the partition and click the **Remove** button. You can select more than one partition by holding down the **Ctrl** key while selecting multiple partitions with the mouse.

A message will be displayed, warning you that you are about to delete all the data deposit copies written on the partition(s). Click **Yes** to remove the partition or **No** to cancel the action.



Warning

Removing a local partition will *delete* all the data in that partition. Therefore, if the partition you remove was associated with a pool for which there are no policy rules defined that guarantee that the system is storing more than one copy, you will be effectively deleting the only copy of the data that is stored in the Circle. A warning message will be displayed if the removal will result in the last copy of a file being deleted permanently.

For information on protecting your data using the storage policy, see [Chapter 15, Storage Policy](#).



Note

The functionality for deactivating and removing *shared* partitions differs from that of local partitions. For more information, see “[Deactivating or Removing a Partition](#)” in [Chapter 10, Shared Disk Pools](#).

Node and Policy Information

Each disk pool has two tabs showing node and policy information.

The **Nodes** tab shows the node the disk pool has been configured on.

The **Policy** tab lists the policies which are using the pool.

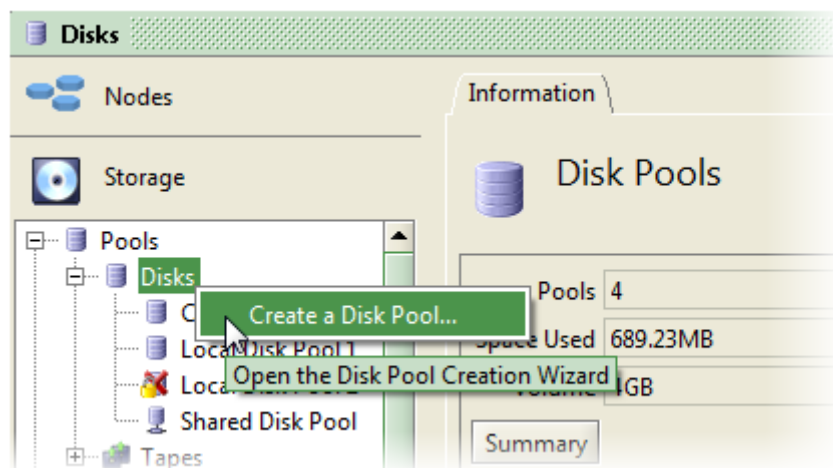
Chapter 10. Shared Disk Pools

Shared disk pools are defined based on shared storage partitions in the Circle. A shared partition can be defined before or during the creation of the shared disk pool. These partitions must be available and recognizable to the servers in the network through a LUN (Logical Unit Number) or as an exported NFS volume. Each partition must then be declared in Active Circle using its volume label. Once a shared disk pool is created and validated, it will be available to all the nodes you have given access to. The pool will be managed through a *Leader Node*, and the shared disk pool element will be attached to this lead node in the **Nodes** view.

Creating a Shared Disk Pool

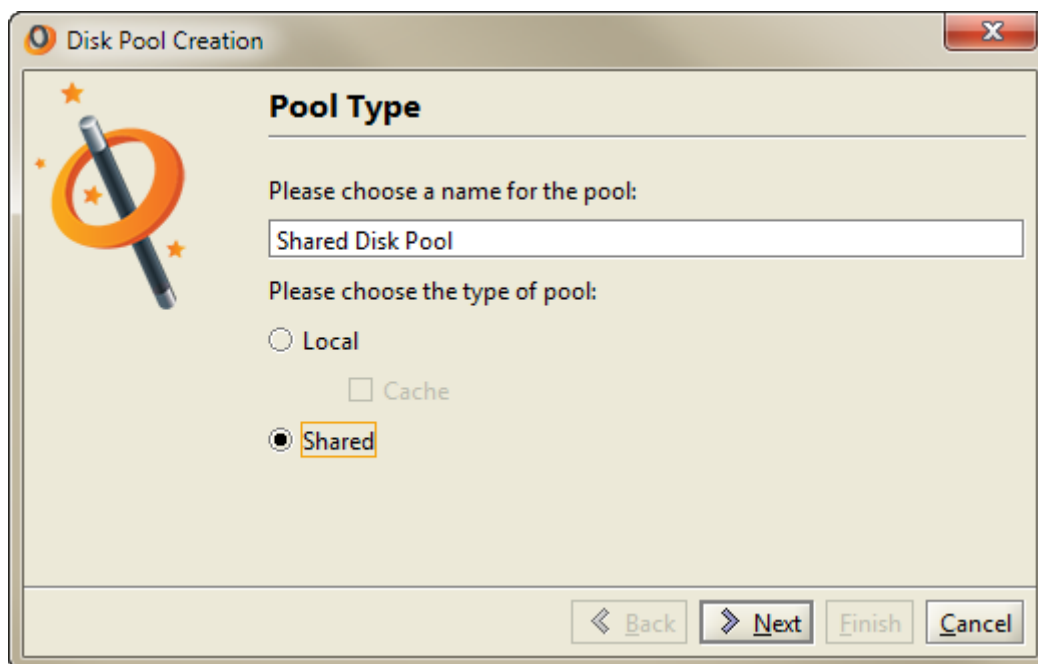
To create a new pool, right-click on either the **Pools** or the **Disks** element at the top of the **Storage** view. Select **Create a Disk Pool** from the context menu. This will start the pool creation wizard.

Figure 10.1. Starting the Pool Creation Wizard



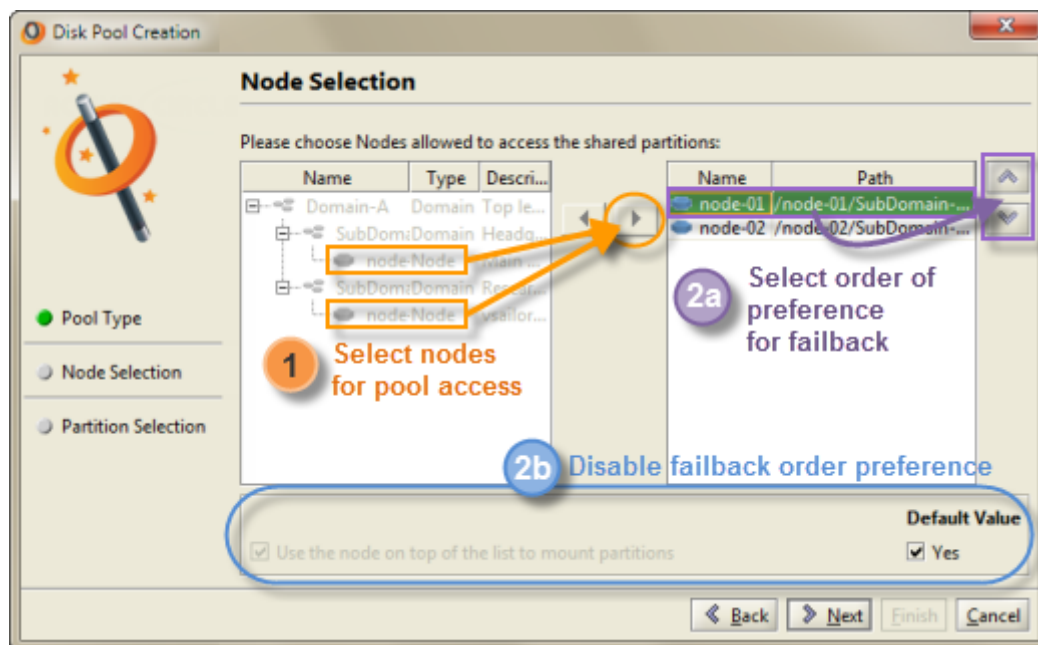
In the first panel, you must enter a name for the pool and choose the type of pool to create. Select **Shared**, and then click on Next.

Figure 10.2. Pool Name and Type



In the next panel, you must select the nodes which will be allowed to access the shared storage partitions for this pool, as well as the node order of preference. Hold down the **Ctrl** key to select multiple nodes from the list of available nodes. Click on **Next** to continue.

Figure 10.3. Selecting the Access Nodes and Order



This is done in two steps:

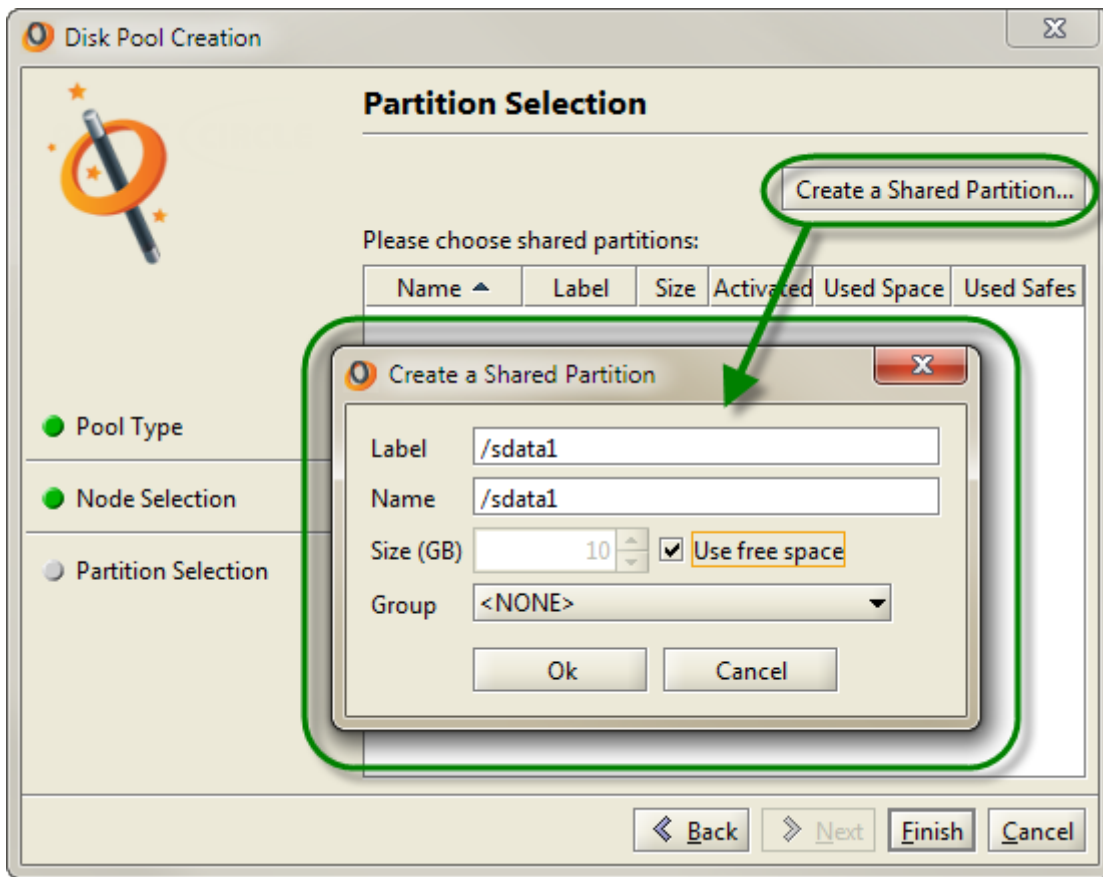
1. The list on the left displays the available nodes in the Circle. Select each node from the list and click on the arrow to add them to the list on the right.
2. Next, you need to specify the failback mode and node order for the pool. There are two options for this:

- **Option 2a:** Use the up and down arrows to the very right of the window to select the order of preference for the nodes. The node at the top of the list will become the *leader node*, on which the partition will be mounted and access to the pool will be managed. The pool will automatically fail back to this node whenever it is online. This the default mode.
- **Option 2b:** If you don't want to use the automatic failback mode, but instead want to keep the first available node as leader as long as it is in service, you need to deselect the greyed-out option at the bottom of the window. To access this field, you must first click to remove the check mark in the box marked "Yes" under **Default Value** (if this box is labeled "No", the default value has been modified at the Circle level, and you don't need to change anything here). Then click to remove the check mark next to the field **Use the node on the top of the list to mount partitions**. You can still set the order of the nodes using the arrow buttons to the right. The order will be respected when a leader re-election needs to take place.

For more information about what these options mean, read the sections on order preference and failback mode under [Managing Nodes in the Shared Disk Pool](#). There you will also find a description of how to modify the node selection, the order of preference and the failback mode after the shared pool has been created. When done, click on **Next** to continue.

The **Partition Selection** panel lists unassigned shared partitions that have been defined in the Circle. If the partition you want is on the list, select it and click on **Finish** to close the wizard and create your pool (and ignore the rest of this section). If the list is empty, or you want to define a new shared partition for this pool, you can click on the **Create a Shared Partition** button. The window that is displayed for entering the partition information is the same (except for the group selection list) whether you define the shared partition from here or from the **Disks** element of the **Storage** view.

Figure 10.4. Defining the Shared Partition



Fill in the fields to identify the shared partition, following these instructions carefully:

Label

The label should be a unique identifier of the partition in the network. It can be a LUN corresponding to the actual volume label of a partition as defined in a SAN environment. The label can also be specified as an NFS mount point in the format `servername:/path/to/partition`. The resource represented by the label *must* exist and be visible to all the nodes that will be included in the shared pool using the partition. **Important:** Be careful to spell the label exactly like it is defined in the file system where it is hosted. If you type an incorrect name, the system will accept it, but the partition will not be available.

Name

This field will be filled in with the label text by default. You may specify a different name if necessary.

Size (GB) / Use free space

Specify the full size of the partition in gigabytes, as the whole volume must be reserved for Active Circle. The simplest way of doing this, is to check the box **Use free space**, to let the system determine how much space is available when the volume is mounted and use all of it for the shared partition. The partition should be empty. If it contains any data, it may be unusable for shared disk pools.

Group

If you have created one or more *Shared Partition Groups*, and you want this partition to be part of a specific group, you can select one here. You can also create groups and move shared partitions to them later. For more information, see [“Shared Partition Groups”](#).

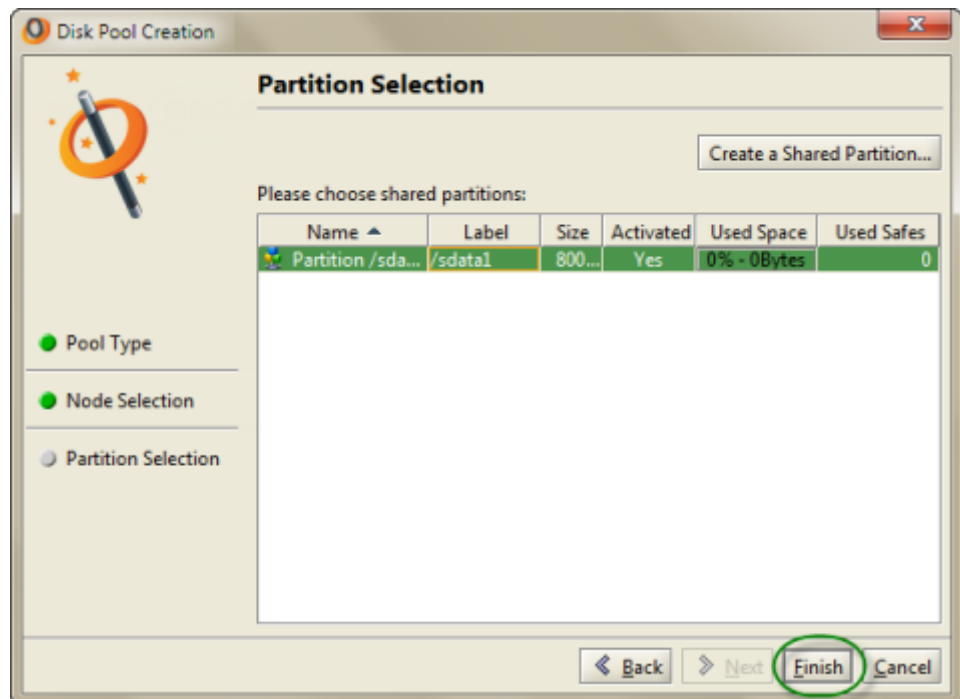


Tip

For more information on *shared partitions*, see [“Partition Types”](#) and [“Adding a Shared Partition”](#) in Chapter 13, *Disk Devices*.

Click on **OK** to create the partition. Then select it from the list and click on **Finish** to create your pool.

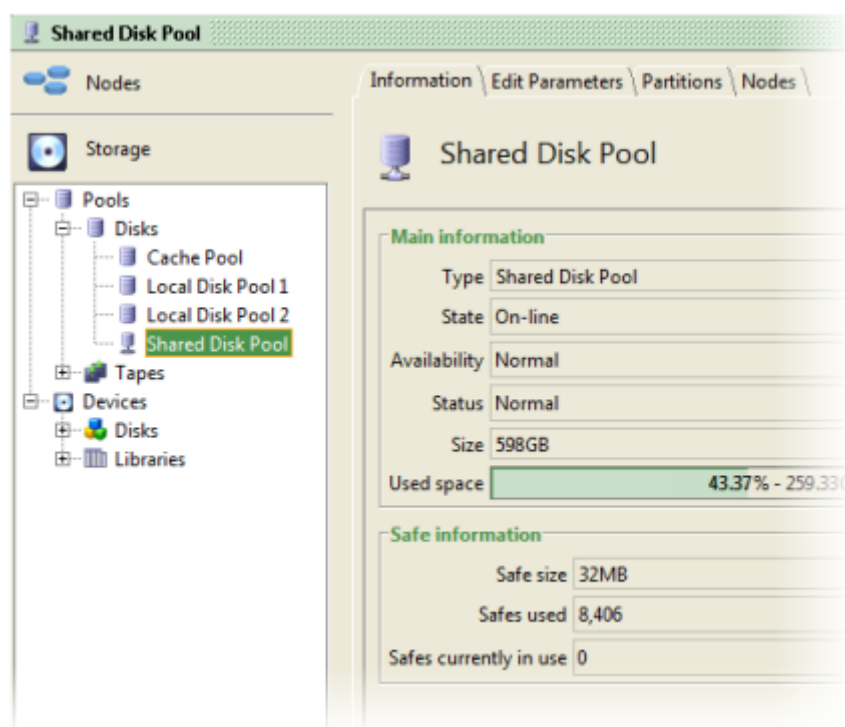
Figure 10.5. Selecting the Shared Partition



Pool Information

The **Information** tab displays all the information about the selected pool.

Figure 10.6. Pool Information tab



Main information

Type	Type of pool (Local or Shared).
State	Current state of the disk pool. Possible values are: On-line, Off-line.
Availability	Current availability of the disk pool. Possible values are: Normal, Read only, Error.
Status	Current status of the disk pool.
Size	Size in GB of the disk pool.
Used space	Amount of disk space currently being used by the disk pool.

Safe information

Safe size	Size in MB of the Active Circle safes (also called packs). Defined during the installation procedure.
Safes used	Number of safes to which data has been written.
Safes currently in use	Number of safes currently open for use.

Changing Pool State and Availability

The settings for **State** and **Availability** are the same for shared disk pools as for local disk pools. Please see:










- [“Disk Pool States”](#)
- [“Disk Pool Availability”](#)

Pool Status Indicators

Visual status indicators help identify the location of any potential issues affecting the disk pools. The indicators are shown as small icons in the views pane, and as larger icons on

the **Information** tab for the element. To find more information when a problem is indicated, please consult the supervision notes (see [Chapter 32, Supervision](#)).

Table 10.1. Shared Disk Pool Status Indicators

Indicator	Meaning
	Shared Disk Pool in normal operation, no problems detected.
	The pool is <i>Read-Only</i> or <i>Full</i> . No more data can be written to the pool.
	The pool is <i>Off-line</i> .
	Initial state (status UNCHECKED). No partition mounted on leader node, and no partitions with error on leader node. Also shown when no partition information is available.
	One or more partitions are correctly mounted on the leader node, but at least one partition has an error.
	All partitions have an error on the leader node (mount failed or partition unavailable).
	Error mounting a partition on the leader node.. A <i>Supervision Note</i> has been issued.
	The indicators can be combined to show multiple problems. In this case, the pool has an error, it is off-line and one or more supervision notes have been issued due to a problem with the pool's partitions.
	The pool is off-line and has been put in maintenance mode. This can happen when no shared partitions are available to the pool.

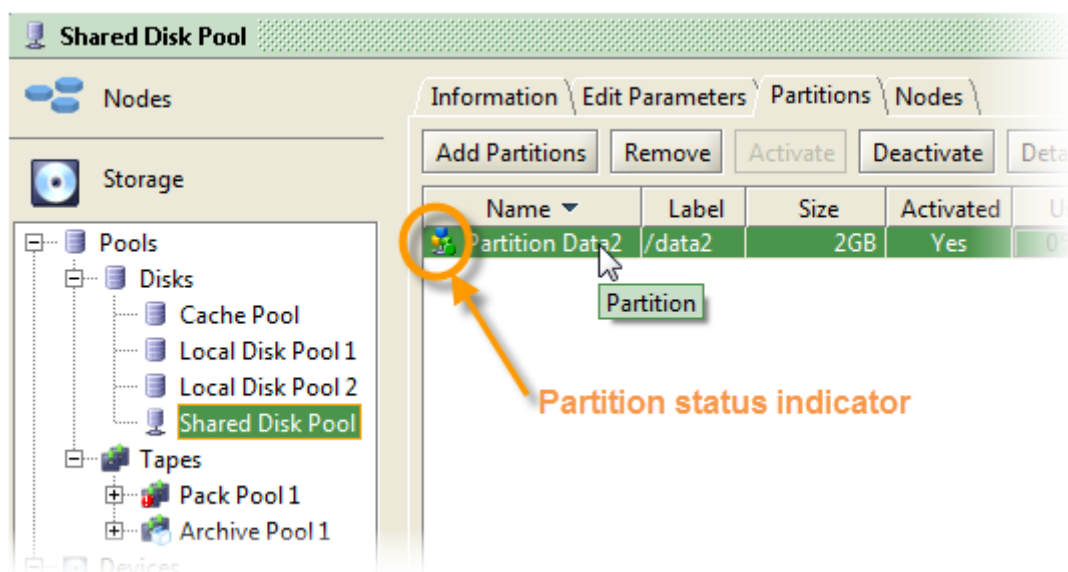
Adding Partitions to a Pool

You can add partitions to an existing shared disk pool. Select a pool in the **Storage** view and click on the **Partitions** tab. Then click on the **Add Partitions** button. In the **Add Shared Partitions** window, you can either select an available partition from the list (if any), or you can click on **Create a Shared Partition** to create a new partition to add. For information on how to fill in the fields when creating a shared partition, see [Defining the Shared Partition](#).

Shared Partition Status

The **Partitions** tab for a shared disk pool displays the list of shared partitions defined for the pool along with information for each partition. The icon to the left of the partition name is a dynamic status indicator which changes to reflect the current status for the partition:

Figure 10.7. Shared Partition Status



The status indicator shows the state of the partition on the current leader node (the node highlighted in grey in the list in the **Nodes** tab for the pool). To display status information for the other nodes in the pool, see [“Partition Access Details”](#). Partition status indicators for the leader node:

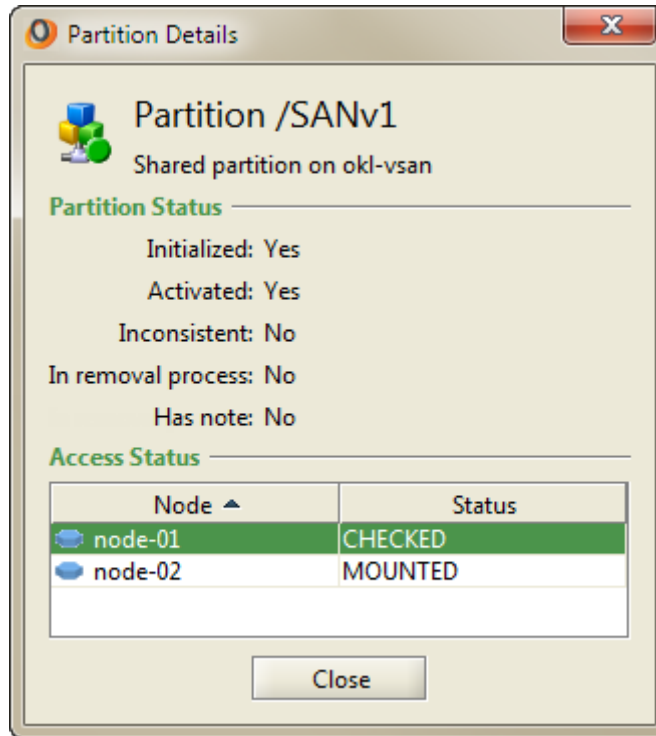
Table 10.2. Shared Partition Status Indicators

Indicator	Meaning
	Normal state. The shared partition is in the MOUNTED state on the leader node, and therefore available for use. For a non-leader node, the normal state is CHECKED.
	The partition is in a transitional state on the leader node, which could be one of the following: UNCHECKED, RESERVING, CHECKING, MOUNTING, UNMOUNTING, RELEASING or PREEMPTING. On a non-leader node, if the system is unable to access or locate the disk volume specified by the partition, the status will remain as UNCHECKED, and a supervision note will be issued.
	The partition has a problem on the leader node. The detailed status can be one of the following: RESERVE_FAILED, CHECK_FAILED, MOUNT_FAILED or UNMOUNT_FAILED.
	The yellow triangle with an exclamation mark indicates that the partition has an error on one of the non-leader nodes of the pool. A non-leader node could be in one of the following states: RESERVE_FAILED, CHECK_FAILED, MOUNT_FAILED, UNMOUNT_FAILED, RELEASE_FAILED, PREEMPT_FAILED, UNAVAILABLE. The green dot in this example indicates that the partition has the MOUNTED status on the leader node.

Partition Access Details

You can verify the detailed partition access status for the nodes in the shared disk pool. Select a partition from the list and click on the **Details** button. The information window shows the status of the partition on each node that has been given access through the shared pool.

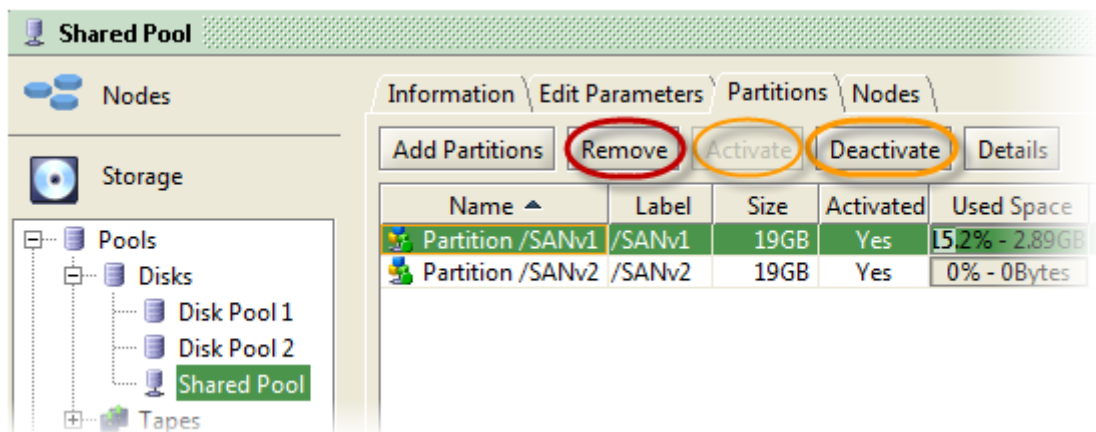
Figure 10.8. Access Status Details Window



Deactivating or Removing a Partition

You can deactivate or remove one or more partitions using the **Partitions** tab of a shared disk pool.

Figure 10.9. Deactivate, Activate and Remove Buttons



Deactivation

Select the partition(s) in the list and click on the **Deactivate** button (hold down the **Ctrl** or **Shift** key to select multiple partitions). *Deactivating a shared partition will cause it to be unmounted*, and as a consequence the data on the partition will not be available for reading or writing. When deactivated, the shared partition will be set to the status **UNAVAILABLE** and its icon will be greyed out. No data will be erased.

If you deactivate all the partitions of a shared disk pool, the pool itself will no longer be available. The availability status will transition from *Normal* to *Full*, then to *Error*. The pool will be taken off-line. Supervision notes will be issued to inform the administrator about the status.

Reactivation

You can make the partition(s) available again by selecting one or more of them and clicking on the **Activate** button. Note that if you reactivate partitions in a pool which is in the state *Off-line* and/or has availability *Error*, you must manually set the pool back online and make the availability normal before the shared disk pool can be used. This is done on the **Edit Parameters** tab of the pool.

Figure 10.10. Setting Shared Pool State to On-line

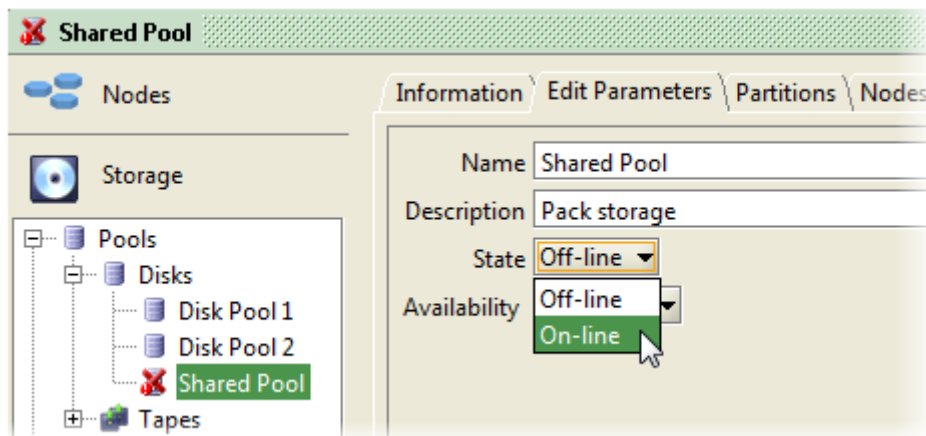
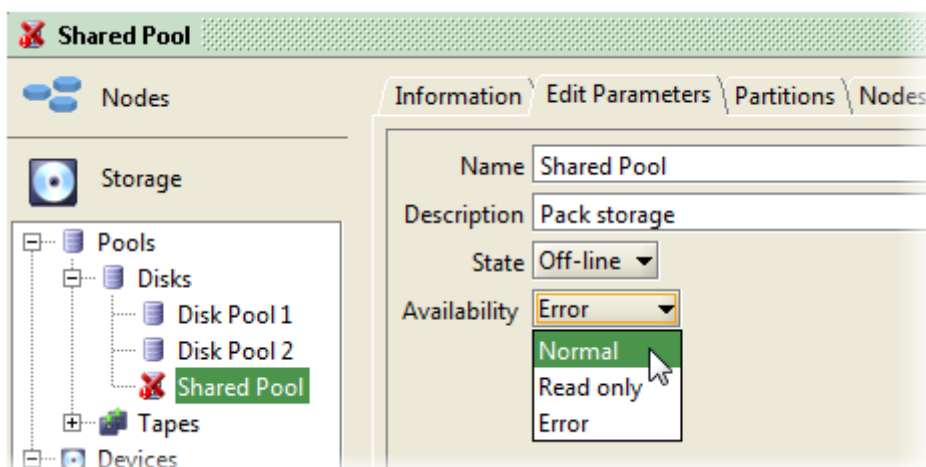


Figure 10.11. Setting Shared Pool Availability to Normal



The system will not update the pool state on partition reactivation automatically, as the administrator may have decided to set and keep the pool off-line independently of the partition status.

Removal



Note

Make sure the partition is currently in *activated* mode before removing it (unless there is an error condition preventing activation). You should also verify that the pool is on-line, as indicated by the **State** field on the **Information** tab for the pool. Partition removal will not be processed correctly if the pool is off-line.

If you remove a partition from a shared disk pool, all the Active Circle data deposits on the partition will be lost. Select one or more partitions and click on the **Remove** button if you wish to do this. You will receive a warning about data being deleted which you will have to confirm. If you click **Yes** to continue, the data will be deleted and the partition will be unmounted. The partition will be removed from the list of partitions for the pool.

Partitions that have been removed from a pool will still be visible in the **Storage** view under the **Devices > Disks** element.

Removing all the partitions of a shared disk pool will set the pool availability status to *Full*, and a supervision note will be issued informing the administrator that there is no more storage space in the pool.



Note


If you remove a partition that is already in a deactivated state (either due to a manual administrator action or automatically as a result of an error), the data on the partition will not be erased. This is because the data on the partition are no longer available to the system when the partition has been deactivated. Although the physical data may not be deleted, the metadata references will be erased, and the partition will no longer be usable in the system in the resulting *inconsistent* state. To reuse such a partition, you must follow the manual procedure described below.

Reusing a shared partition that has been removed

A shared partition that has been removed from a pool and had its contents erased automatically in a normal way (without being deactivated first) can be reused simply by adding it to the same or another shared disk pool. To do so, use the **Add Partition** button on the **Partitions** tab for the pool and select the partition from the list that is displayed.

If, however, the partition was removed after it had first been deactivated, it will be in an inconsistent state and cannot be reused without manual preparation. If you try to add such a partition to a pool, the system will not be able to mount it. The file system check will fail, and the partition will remain unavailable. A supervision note will be issued informing the administrator that the partition is inconsistent. The detailed partition status will look like this:

Figure 10.12. Partition Status Inconsistent



To prepare an inconsistent partition for reuse by a pool, follow these steps:

1. Outside of the Administration tool:
 - a. Mount the partition manually using the Linux **mount** command.
 - b. Delete the contents of the partition.
 - c. Unmount the partition (Linux **umount** command).
2. In the Administration tool, go to the pool, select the partition and **Activate** it.

The system should now be able to mount the partition and make the storage space available to the pool.



Note

If all the partitions of a share pool are inconsistent, the state and availability of the pool will be in *Error/off-line*. The pool will stay like this even after the partitions have been reactivated. To resolve this, you must manually update the state and availability of the pool, as described [above](#).

Managing Nodes in the Shared Disk Pool


The **Nodes** tab displays the list of nodes in the shared disk pool and their status. The system will by default fail back on the node which is ranked highest on this list. You can define the order of preference of the nodes to make sure the system always tries to fail back on the nodes you have placed at the top of the list, even when another node is already running as the leader. This failback behavior can be disabled if you do not want the system to perform automatic re-election based on a set order of preference. This will be described in the following sections.

Since the shared partition(s) of the pool are mounted on the leader node, all operations on the partitions pass through this node. If the leader node should become unavailable or lose its connection to the shared storage (for example as a result of a host bus adapter malfunction) ,

the system will start a process to elect a new leader node by trying to mount the partitions on another node according to the failback setting of the system. The connectivity detection for the pool will occur within an interval of maximum 5 minutes.

The line with the current leader node will always be highlighted in grey in the list:

Figure 10.13. Nodes tab for a Shared Disk Pool



In the table, the **State** column indicates the state of the node server itself. *In Service* means that the server and the pool services are running. *Off* indicates that Active Circle is not active on the server. *Start in Progress* means the server is starting. *On* means that the server is started, but some pool activation processes are not yet completed. *In maintenance* is the status indicated during a reset or validation of the pool, during which the shared partitions are not available.



Note



An election of a new leader node automatically takes place if the current shared storage leader has been silent for a certain time period. By default, this period is 50 seconds. After this, the shared pool will switch to another node. You can change the maximum duration before re-election takes place, using the *circle property* `diskPool.clusterSilenceDuration`. The value is specified in seconds, and it must be between 5 and 120. It applies immediately, no restart required. For more information on how to set values for circle properties, see [“Circle Properties”](#).

Node Status Indicators

The icon to the left of the node name in the list is a dynamic icon that changes to represent the current state of the node in relation to the shared partitions in the pool.


Table 10.3. Node Status Indicators

Indicator	Meaning
	All the partitions in the pool are in the normal, functioning state on this node. For a leader node, the status is MOUNTED. For a non-leader node, the status is CHECKED.
	At least one partition has a normal status on this node, and at least one partition has a problem (but not necessarily an error).

Indicator	Meaning
	No partitions are in normal working status on the node, but not all the partitions have an error. This can mean that the status is UNCHECKED, which can be corrected/verified by validating the configuration. If a secondary node has the status UNCHECKED, it will have no impact on the functionality of the pool as long as the leader node is in the normal (MOUNTED) state.
	All the partitions of the pool have an error (status *_FAILED or UNAVAILABLE) on the node.

To see more detailed partition access information for a particular node, select the node from the list and click on the **Details** button.

Figure 10.14. Node Access Status Details

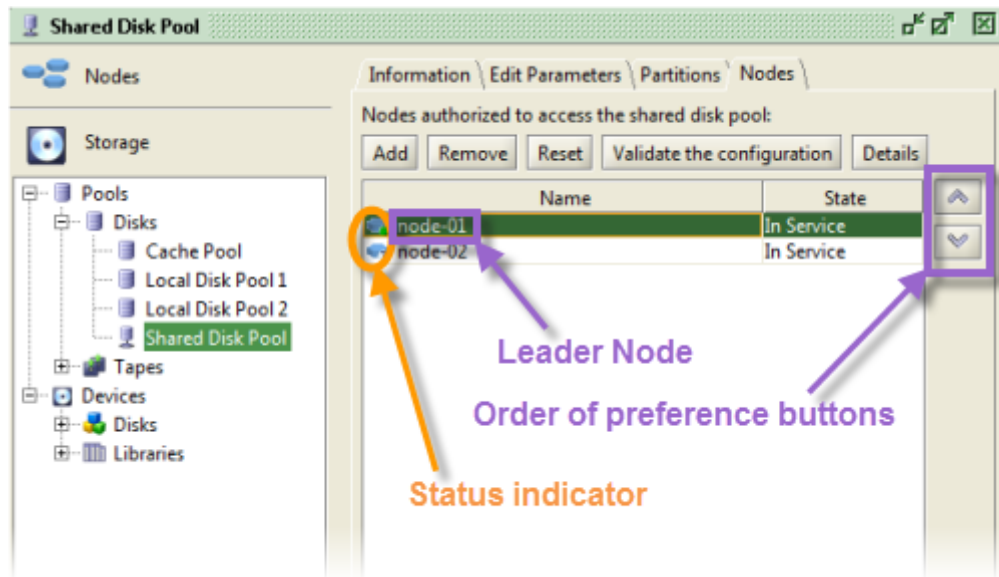


Changing the Order of Preference

The **Nodes** tab allows you to change the order of preference of the nodes in the shared disk pool. The default failback behavior is that the order of precedence set by the administrator in this window will be used. The node at the top of the list is the preferred leader node. As long as this node is operating normally, it will remain the leader which all data requests go through. If this node goes offline, the second node on the list will become the leader node (after re-election) and take over all the data requests for the storage in the shared disk pool. If the second node is offline, the system will instead try to mount the partitions on the next node in the list, and so on. If a node which is ranked higher than the current leader node comes back online after a failover, a re-election will take place and the system will fail back to the node with the highest ranking, even if the previous leader is running normally (it will lose the re-election because it is placed lower on the list).

To change the order, select the node you want to move up or down, and use the arrows to the right of the list to change the rank of the node in the list. Click on **Cancel** to undo the move. Click on **Apply** to confirm so that the changes take effect. If the arrow buttons are disabled (greyed out), it means that the order preference failback mode has been disabled. For more information on this, see the [next section](#).

Figure 10.15. Indicators and order of preference

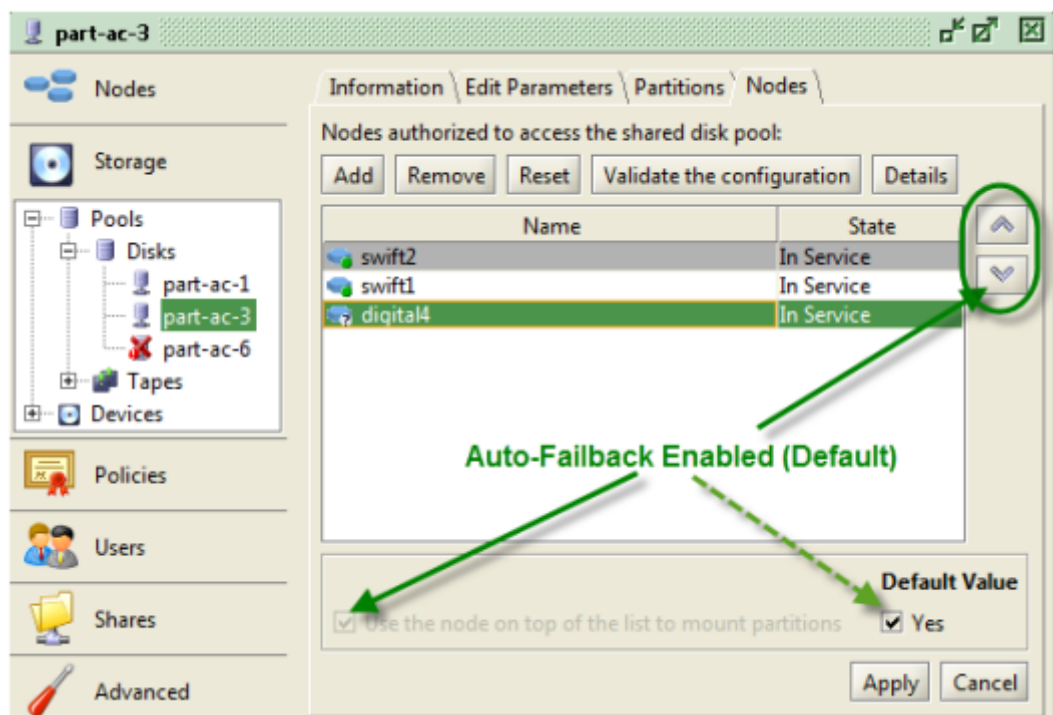


For the changes to be applied to the pool, and to trigger a re-election of the leader, you need to either **Reset** or **Validate** the pool configuration, as described in the following sections.

Disabling the Order of Preference for Failback

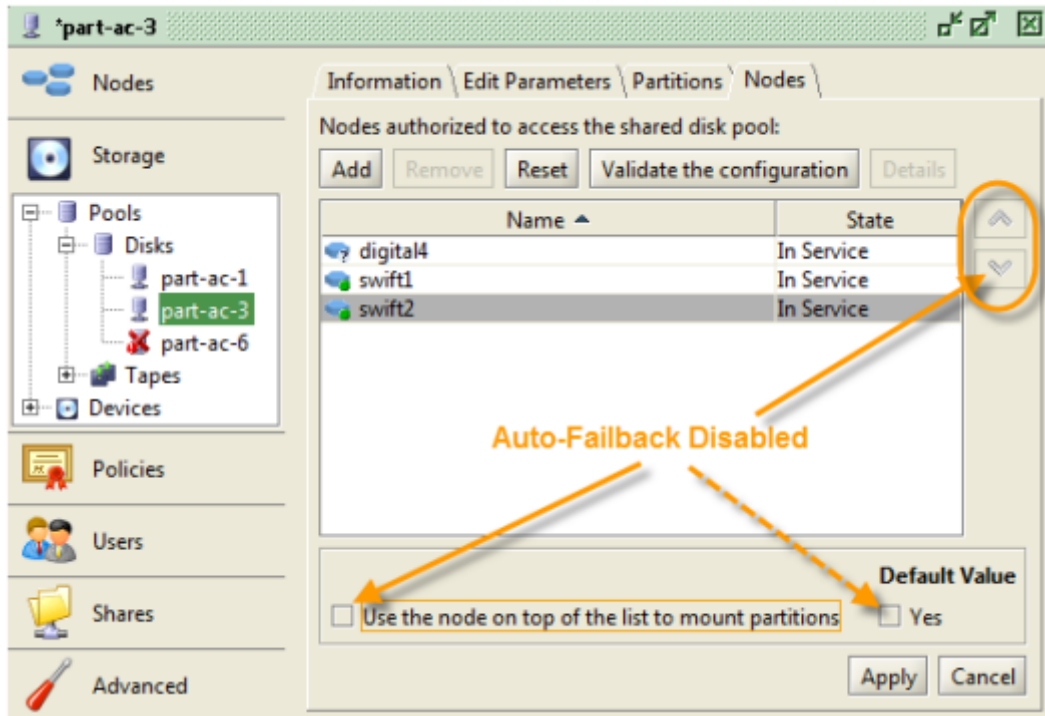
The mode where the pool automatically fails back on a node according to a set order of preference can be disabled at the level of the pool and at the level of the Circle as a whole. For the pool, this setting can be modified during creation of the pool (see [the pool creation wizard](#)) or by using the settings on the **Node** tab for the shared disk pool:

Figure 10.16. Failback Preference – Default Setting



To change the setting, you must first click on the "Yes" check box under **Default Value** to *remove* the check mark. This will make the field labeled **Use the node on the top of the list to mount partitions** available for deselection. To disable automatic failback, you need to click in this box to *remove* this check mark as well. The window should then look like this:

Figure 10.17. Failback Preference Order Disabled



In this mode, there will be no automatic re-election of a leader just because there is a higher-ranked node available. As long as the current leader node runs normally, it will remain the leader. The order of the nodes will still be respected when a re-election takes place. In most cases, the node that starts first will remain the leader when the other nodes come online. Over time, this would normally mean that the leader node responsibility will change less often, resulting in less frequent unmounting and mounting of the partition(s) of the shared pool.

The system will issue a supervision note as information each time a new leader node is elected.



Note

If you want to disable this setting for the whole Circle so that automatic failback according to the node order is disabled for all the shared pools, you need to modify the *circle property* `diskPool.autoFailback` (previously `storage.autoFailback`). Set the value to **false** (the default value is **true**). This will change the **Default Value** field to "No" on all the **Node** tabs for all the shared disk pools, as well as in the pool creation wizard. For more information on how to set values for circle properties, see ["Circle Properties"](#).

Adding and Removing Nodes


You can add and remove nodes in the shared disk pool dynamically while the nodes are running, without interrupting access to the pool on the other nodes. In the **Nodes** tab, click on the **Add** button and select an available node from the list in the **Add Nodes** window. Click **OK** to close the window, and click **Apply** at the bottom right of the **Nodes** tab to confirm the addition.

To remove a node, select the node from the list and click on **Remove**. The node will be removed from the list immediately, but it will not be removed from the pool until you click on the **Apply** button. If you change your mind before you click **Apply**, you can click on **Cancel** to make the node reappear in the list of nodes.

Resetting the Pool

If you change the preferred leader node, you need to reset the shared disk pool. This action tells the system to unmount the shared partitions from the current leader node, perform leader node election and mount the partitions on the new leader node. The data in the pool will be unavailable during the operation. To reset the pool, click on the **Reset** button on the **Nodes** tab. A warning will be displayed. Click on **Yes** if you want to proceed.

Figure 10.18. Reset Pool

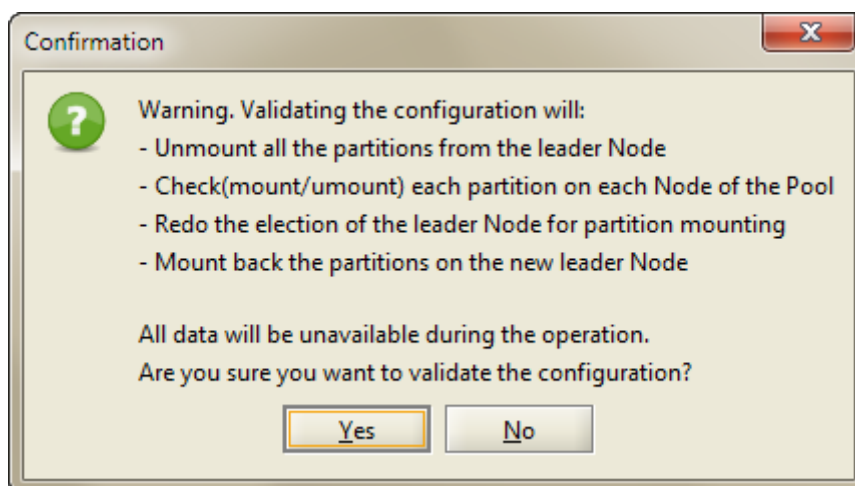


Validating the Pool Configuration

A validation of the pool configuration will perform the same operations as a reset, but in addition it will perform a check for all the partitions on all the nodes in the pool. This means the system will try to mount and unmount all the partitions on each node of the pool. This is, in fact, a validation of the failover/failback capability of the shared disk pool.

After the initial configuration of the pool, if the partitions are successfully mounted on the node that has been elected leader, the system will not try to mount on any of the other nodes. It is therefore recommended to perform this validation. By running this test, the system verifies that the shared partitions can be mounted on each node in the pool, ensuring that mounting on a new leader node will be possible in a failover situation if the current leader fails. The data in the pool will be unavailable during the operation. To validate the pool, click on the **Validate the Configuration** button on the **Nodes** tab. A warning will be displayed. Click on **Yes** if you want to proceed.

Figure 10.19. Validate the Configuration



During the validation, the system will perform the following sequence of tasks for each partition on all the nodes of the shared pool:

Table 10.4. Validation sequence

Operation	State if successful
RESERVING* > CHECKING** > MOUNTING	MOUNTED (intermediate)
> UNMOUNTING > RELEASING*	CHECKED (final)

* performed by SCSI commands

** performed by the **fsck** Linux command

In addition, for the leader node, an initial unmounting will be performed at the beginning of the process as well as a final mounting procedure after all the other nodes have been validated and leader election has taken place.

If the validation fails, a supervision note will be issued. The state of the pool will not be changed.

Chapter 11. Tape Pools

In order to use the tapes in a library for storage or archiving, you must create a *Tape Pool* dedicated to this purpose. A tape pool is a group of tapes that together form a single storage resource in Active Circle. Once you have created a pool, you can assign tapes to it. The storage capacity of a tape pool is the combined sum of the storage capacity of the tapes in that pool. This chapter describes the types of tape pools in the system, how to [create tape pools](#) and how to manage the pools and their tapes.

There are two main types of tape pools that the administrator can create:

- **Archive Tape Pool** — A group of tapes used for archiving in TAR or LTFS format.
- **Pack Tape Pool** — A group of tapes storing data in the Active Circle *pack* format.

In addition, Active Circle automatically organizes all tapes into categories by creating [Default Tape Pools](#).

Tape pools are managed by using the **Tapes** element under the **Pools** category in the **Storage** view. All the tape pools created in Active Circle will be attached to this sub-category. The **Tapes** element has two general tabs; the **Information** tab, which lists all tape pools with summary information, and the **Tapes** tab, which lists all the tapes of all categories and pools attached to the libraries in the Circle.

Figure 11.1. Tape Pools and Tapes tab

The screenshot displays the 'Tapes' management interface. On the left, a tree view shows the hierarchy: **Storage** > **Pools** > **Tapes**. Under **Tapes**, there are sub-items for 'pack', 'tar2-5', 'tar3-0', etc. The main panel shows the 'Tapes' tab with a table of tape pools. The table has columns: Name, Bar Code, Type, Size, Used Space, and Data Size. The table lists various tape pools, including 'tar2-5', 'tar3-0', 'tar-1', 'tar2-4', 'tar2-3', 'pack-2', 'pack-1', 'tar-3', 'tar2', 'tar2-3', 'tar2-4', 'tar2-5', 'tar3', 'tar3-0', 'Blank Tapes-15', 'Unregistered Tapes-9', and 'Unregistered Tapes-17'. A purple circle highlights the 'Default Tape Pools' section, and an orange circle highlights the 'User-defined Tape Pools' section. A purple arrow points from the 'Default Tape Pools' section to the 'Cleaning Tapes' tab in the top navigation bar.

Name	Bar Code	Type	Size	Used Space	Data Size
tar2-5	NKG650L2	LTO-2	367.53GB	25.75% - 94.65GB	100.95GB
tar3-0	NKG644L2	LTO-2	367.53GB	74.68% - 274.47GB	292.05GB
tar-1	NKG643L2	LTO-2	367.53GB	90.45% - 332.44GB	360.01GB
tar2-4	NFC216L2	LTO-2	181.26GB	100% - 181.26GB	191.1GB
tar2-3	NKG640L2	LTO-2	367.53GB	63.52% - 233.45GB	0Bytes
pack-2	NFC220L2	LTO-2	367.53GB	45.67% - 167.85GB	123.97GB
pack-1	NKG647L2	LTO-2	181.26GB	36.15% - 65.53GB	37.18GB
tar-3	AAT744L3	LTO-3	400GB	0% - 0Bytes	0Bytes
Blank Tapes-15	AAT742L3	LTO-3			
Unregistered Tapes-9	CLN819L1	Generi...			
Unregistered Tapes-17	NKG656L2	Generi...			



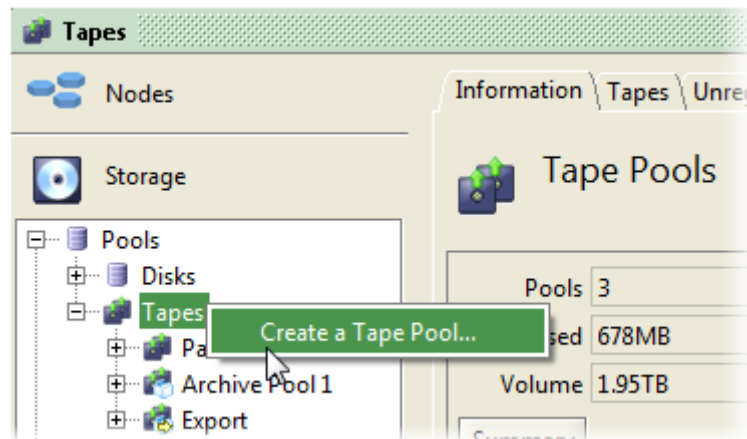
Note

The number of tape pools in the system is limited to 100 by default. You can increase or decrease this limit by modifying the value of the [circle property](#) `media.maxPoolCount`.

Creating and Activating a Tape Pool

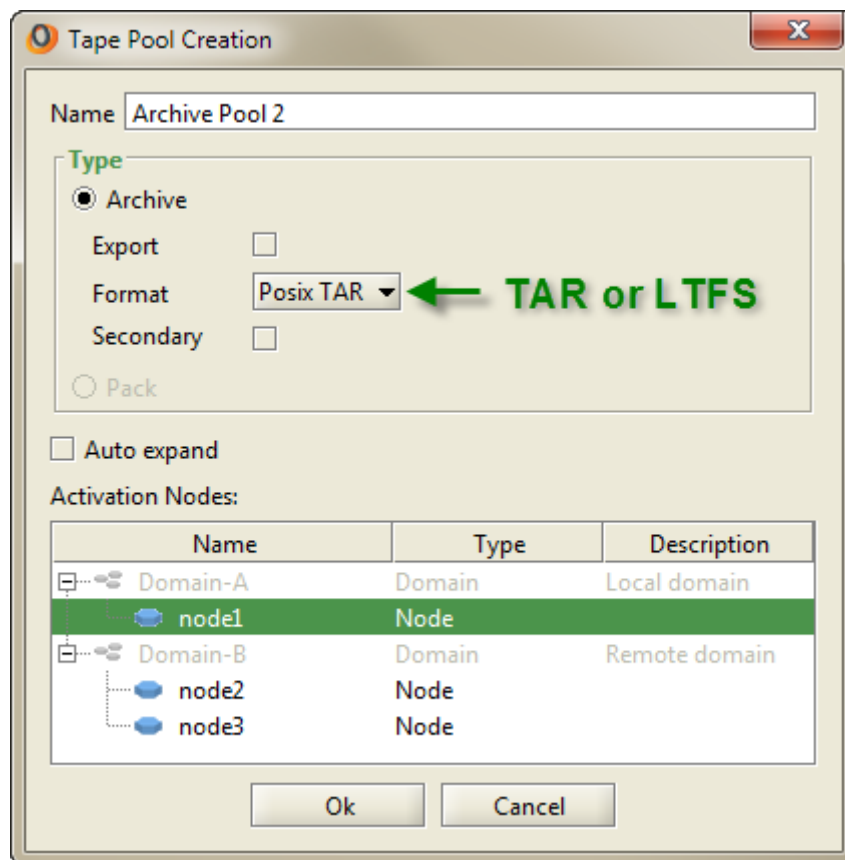
To create a new tape pool, right-click on either the **Pools** or the **Tapes** element in the Storage View and select **Create a Tape Pool**. Alternatively, you can use the **Create** button on the **Information** tab for the **Tapes** element.

Figure 11.2. Create a Tape Pool menu item



The **Tape Pool Creation** dialog box appears.

Figure 11.3. Tape Pool Creation dialog



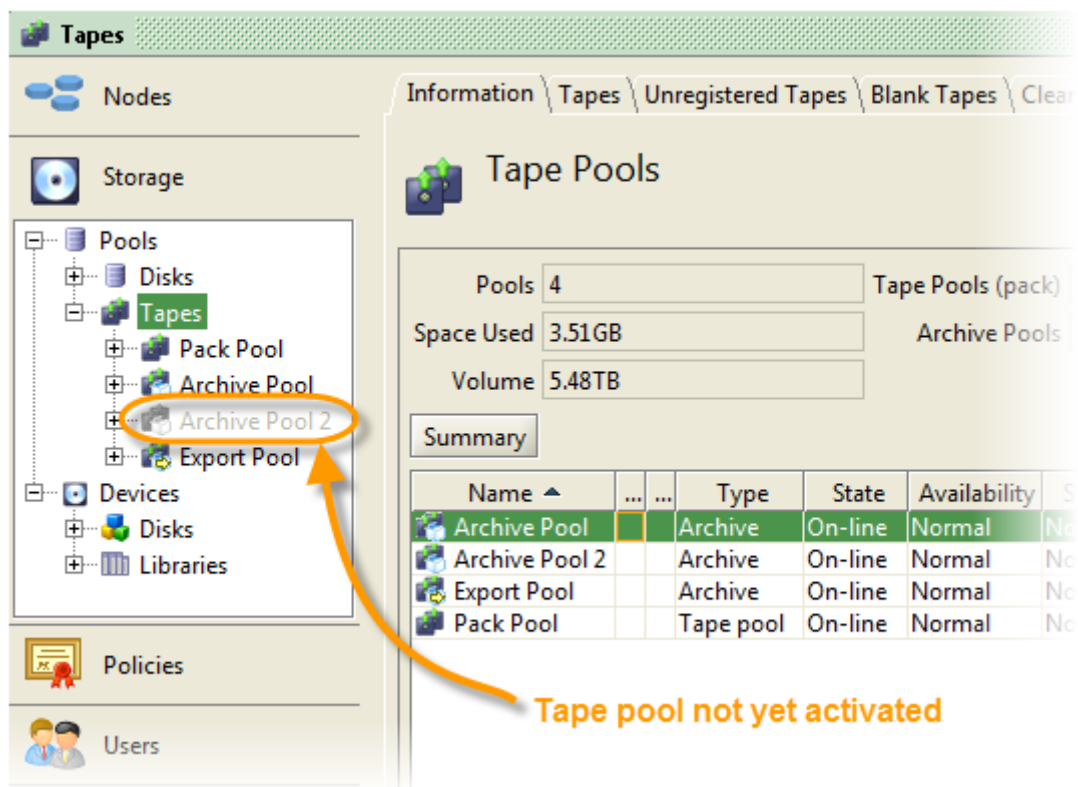
You must enter a name to identify the new tape pool and then choose the type of pool:

Chapter 11. Tape Pools

Archive	A tape pool that will be used for archiving data on tapes. The data stored in this tape pool can be formatted using TAR or LTFS.
Export	You can create a special pool for <i>exporting</i> Posix TAR or LTFS archives from the system. If you are creating an export pool, click the Export check box. This pool cannot be used for regular archiving, so do not choose Export if you want to create a pool to use for archiving of data within the system. You cannot change the type of the pool once it is created. For more information, see the “Export Pool” section.
Format	Choose either Posix TAR format or LTFS format, depending on the type of tape format you plan to use with this tape pool. Note that if you plan to use LTFS, the minimum level supported for tapes and tape drives is LTO-5.
Secondary	Select this option if you want to define the tape pool as a secondary location for reading data from tape when copies of the data are stored in more than one tape pool. This can be useful if the tapes in this pool will be stored offline or at another location than the tape library, so that tape reading operations will use these tapes only as a secondary option in case the primary tape pool for the data is not available.
Pack	A tape pool that will be used as pack storage. The storage format in this type of pool is the same as in Active Circle disk pools. This type of archiving should only be used in special cases, and the option is therefore greyed out by default. If you want to enable creation of pack tape pools, you need to change the value of the circle property <code>media.packEnabled</code> to "true".
Auto Expand	Select whether or not to use the Auto expand feature. When checked, Auto expand will monitor the amount of data written to the tape pool and will automatically add Blank Tapes to the pool if Active Circle detects that the free space is running low. When using this feature, make sure you assign some tapes to the default pool Blank Tapes .
Activation Nodes	A tape pool cannot be used unless it is activated on one or more nodes. You can select the activation nodes in this dialog, or you can do it later as described below. Use the Ctrl key to select more than one node from the list.

When you are done, click **OK** to close the dialog box and create the tape pool. If you did not select any activation nodes, the new pool will appear greyed out in the **Storage** view under the **Tapes** element until you **Activate** the pool on one or more nodes.

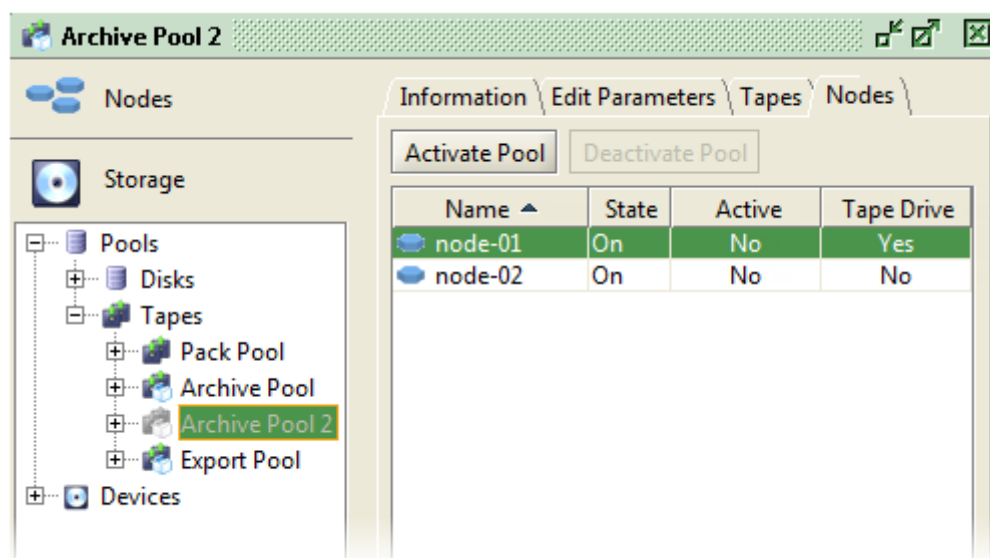
Figure 11.4. Tape Pool Information



Activating a Tape Pool

Before you can use a newly created tape pool, you need to activate it. Select the pool element in the **Storage** view, and click on the **Nodes** tab. The list shows all the available nodes, but only the nodes connected to a tape drive will be eligible for activation. Select a node which has a **Yes** in the **Tape Drive** column (you can select more than one node by holding down the **Ctrl** key while selecting). Then click on the **Activate Pool** button to activate the pool on that node.

Figure 11.5. Activating the Pool on a Node



The pool will be activated instantly, and the **State** column for the node will change from *On* to *In Service*. The pool element will no longer be greyed out in the **Storage** view.



Note

For a pool using a *shared tape library*, the pool must be activated on all the nodes that have access to the library.

The columns in the table shows the following information:

State	Current status of the node. Possible values: In Service, On, Starting, Off.
Active	'Yes' when the tape pool is active for the node. Possible values: Yes or No.
Tape Drive	Presence of a tape drive connected to the node. Possible values: Yes or No.

The pre-defined (default) tape pools do not require activation.



Note

Activating a tape pool for a node here has the same effect as activating the pool in the **Tape Pools** tab for the node in the **Nodes** view.

Default Tape Pools

Active Circle organizes the tapes in the libraries attached to the system into tape pools according to how they are currently being used. Unlike the user-defined pools, these system-defined tape pools (or categories) are not listed under the **Pools** → **Tapes** element in the **Storage** view. Instead, each default pool type is represented by a category tab in the main information window when clicking on the **Tapes** element. Each tab contains information on the number of tapes and a list of the tapes in that category, as well as buttons used to [manage the tapes](#).

Unregistered Tapes

When you attach a supported tape library to a node, Active Circle will recognize the device and will place all of its tapes into the **Unregistered Tapes** pool. It then becomes the responsibility of the system administrator to assign them to the appropriate tape pool.

The **Unregistered Tapes** tab lists all of the cartridges in all the tape libraries that are attached to the nodes in your Circle. It is from here that you can assign the tapes to the appropriate tape pools according to how they are to be used in the system.

Blank Tapes

The pre-defined tape pool **Blank Tapes** serves as a recycling bin for tapes. If, for example, the contents of a tape have been archived to off-line storage and is no longer needed on-line, you can move it to the **Blank Tapes** pool. This will delete the contents of the tape, thereby preparing it for re-assignment to another tape pool.

If **Auto expand** is enabled for a tape pool, it will look for tapes among the blank tapes to add to the pool when it is running low on storage space. It is therefore necessary to have some

tapes in the **Blank Tapes** pool at all times. When the **Auto Expand** feature is being used, the system will issue a supervision note when it considers that this pool is running low on space (only has a few tapes left).



Note

The system will not allow a tape that is loaded in a tape drive to be recycled into the **Blank Tapes** pool. To remove the tape, it must first be unloaded from the drive. Likewise, if a blank tape is loaded in a tape drive, it cannot be assigned to a tape pool using **Auto-expand**. The auto expand functionality will only look for tapes in the **Blank Tapes** pool that are *not* loaded in a tape drive.

Cleaning Tapes

The **Cleaning Tapes** pool is used for keeping track of the cartridges used for cleaning the tape drives. It is not possible to initiate a cleaning operation from here. For information about cleaning operations for a tape drive, please refer to [“Cleaning a Tape Drive”](#).

In order for Active Circle to be able to clean a tape drive, three basic conditions must be met:

- A cleaning tape must be present in the library.
- The cleaning tape must be assigned to the **Cleaning Tapes** pool. For information on assigning tapes to pools, please see [“Managing Tapes in Tape Pools”](#).
- The tape must be in good condition. This means that, depending on the tape format, it should not have been used more than about 20 times.

There are two pieces of information relating to cleaning tapes that are different from other tape categories:

Last Use

For normal tapes, which are used for data storage, this field is **Last writing date**.

Elapsed cleaning cycles

The number of times the cleaning tape has been used to clean a tape drive.

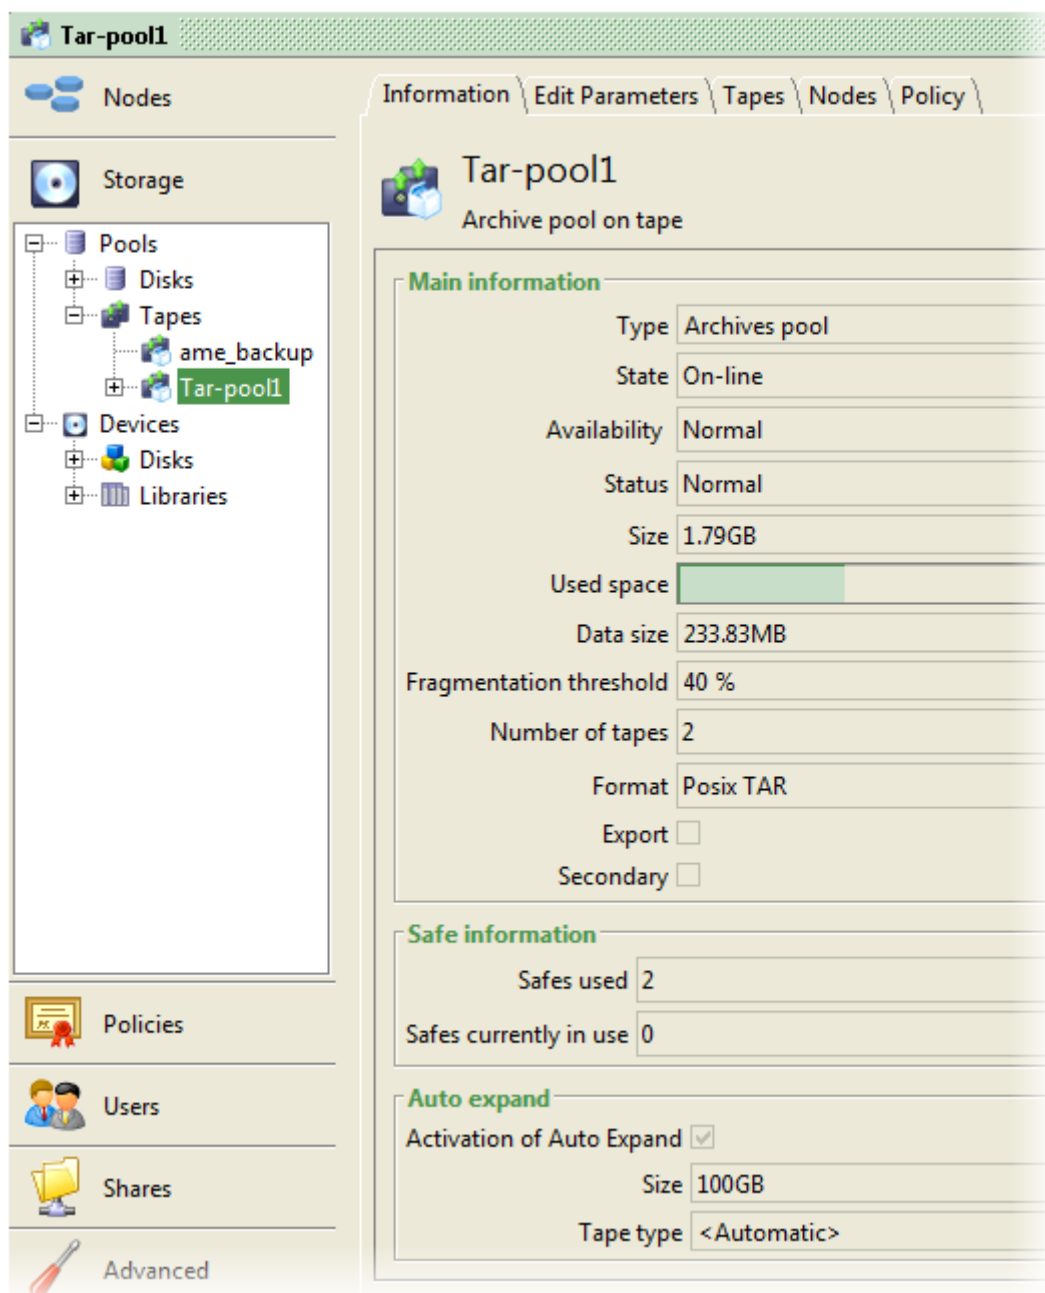
Excluded Tapes

The **Excluded Tapes** pool provides you with a means of ensuring that certain tapes will *not* be used by the Active Circle system for storing or archiving data. This would be the case if, for example, a subset of tapes in a library were being used by another application. By placing those tapes in the **Excluded Tapes** pool, you can be sure that Active Circle will not be writing data to them.

Tape Pool Parameters

When you click on a tape pool item in the **Storage** view, the **Information** tab for the pool presents an overview of the tape pool parameters.

Figure 11.6. Pool – Information tab



Main Information

Type	The type of tape pool: Pack or Archive.
State	Can be either On-line or Off-line .
Availability	Can be either Normal or Read-only .
Status	Shows the status of the pool.
Size	The overall storage capacity of the tape pool (combined size of the tapes assigned to the pool).
Used space	Number of bytes written to tape after possible tape drive hardware compression. This field shows the space the data occupies on the tape after compression. The bar and percentage figure

	represents the share used in relation to the medium's maximum capacity and its remaining capacity as measured by the drive. Compression is performed by the tape drive, and the compression ratio depends on the type of drive.
Data size	The Data size field indicates the actual size of the data on the tape, before compression. This is the size of the data if retrieved from the tape to disk.
Fragmentation threshold	The Fragmentation threshold is the percentage at which a tape in the tape pool is considered fragmented and therefore in need of defragmentation. When this threshold is reached for any given tape in the tape pool, a supervision note is generated that will allow an operator to initiate a defragmentation operation (see Defragmentation for more information). This number expresses the relation between the useful volume of data on the tape and the total volume of physically written data. The default threshold is 40%. It can be modified on the Edit Parameters tab. This functionality does not apply to export pools.
Number of tapes	The number of tapes currently assigned to the tape pool.
Format	The data format on the tapes. Can be "Posix TAR" or "LTFS".
Export	Indicates whether the pool is a pool dedicated to exporting data from the configuration.
Secondary	Indicates whether the pool is treated as a secondary location for tape read requests. During optimized reads where two copies of the data exist on tape, a secondary pool will only be used if the pool containing the other (primary) copy is not available.

Archive File information (Safe Information)

This part shows information about how many archive files have been created in the pool.

Auto Expand

This part shows the auto-expansion settings of the pool. When checked, **Auto expand** will monitor the amount of data written to the tape pool and will automatically add **Blank Tapes** to the pool if Active Circle detects that the free space is running low. When using this feature, make sure you assign some tapes to the default pool **Blank Tapes**. When the auto expand feature is activated for a pool, the administrator does not need to assign tapes directly to the pool.

Activation of Auto expand	Indicates if auto-expand is activated.
Space remaining in the pool	Amount of space remaining in the current Pool.
Tape type	The type of tapes that should be added to the pool. In the Edit Parameters tab, you can choose <Automatic> or select a specific type. The available

types depend on your tape drive configuration. Possible types are:

AIT-1	LTO-3	SDLT-I
AIT-2	LTO-4	SDLT-II
DLT-III	LTO-5	VXA-1
LTO-1	LTO-6	VXA-2
LTO-2	SAIT	



Note

If a blank tape is loaded in a tape drive, it cannot be assigned to a tape pool using **Auto-expand**. The auto expand functionality will only look for tapes in the **Blank Tapes** pool that are *not* loaded in a tape drive.

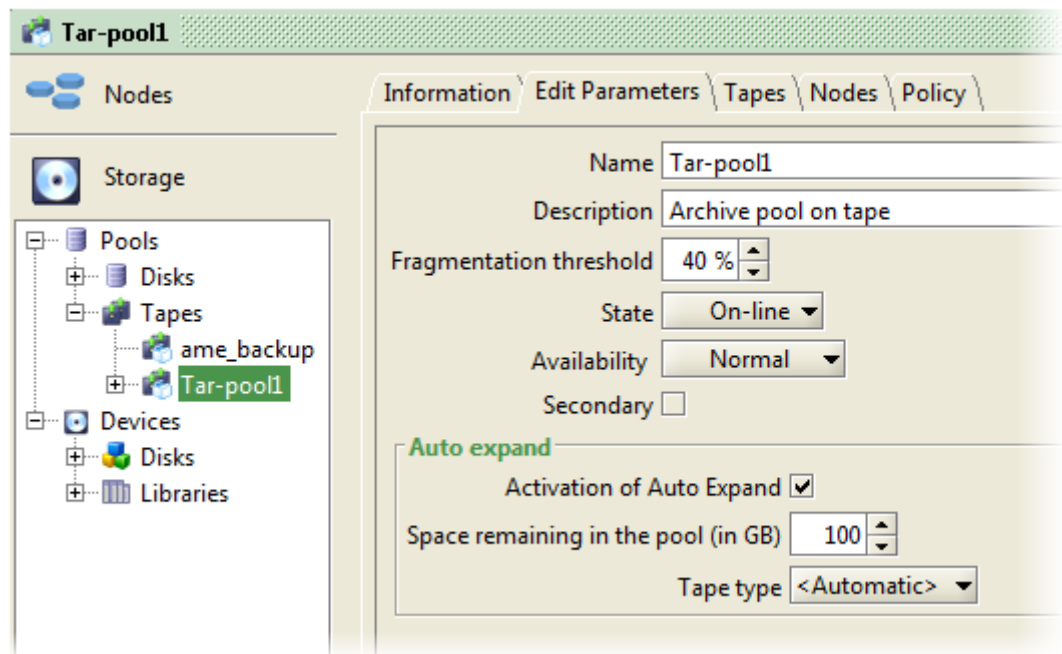
Policy Information

Each tape pool has a separate tab showing policy information. The **Policy** tab lists the policies which are using the pool.

Edit Parameters

The **Edit Parameters** tab presents the pool tape settings that can be modified.

Figure 11.7. Pool – Edit Parameters tab



In addition to modifying several of the parameters described above, you may enter a new **Name** and a **Description** for the pool here.

You can modify the setting for the **State**, **Availability**, **Fragmentation threshold**, **Secondary** and **Auto expand** options.

Tape Parameters

Once a tape (or medium) has been added to a tape pool, you can select it to view its characteristics. Click on the tape under the tape pool item it belongs to in the **Storage** view.

Figure 11.8. Tape in a Pool – Information tab

The screenshot shows a web interface for a tape pool. At the top, there are three tabs: 'Information' (selected), 'Edit Parameters', and 'Advanced'. Below the tabs, there is a header section with a tape icon and the text 'TAR 1-0' and 'Magnetic Tape'. The main content area contains a list of fields with their values:

Bar code	G03003L4
Type	LTO-4
Location	Library 0 - Slot 1026
First assignment	April 30, 2013 4:51:24 PM
Last mounting date	
Last writing date	June 6, 2013 12:02:21 AM
Serial number	G03003L4_1367323635
External location code	
External identifier	
Number of mounts	0
Full	<input type="checkbox"/>


Below this table, there is a section labeled 'Volume' with a 'Size' field showing '800GB'.

The **Information** tab for a tape that has been assigned to a pool contains detailed information about the tape, including:

Bar code	External identifier
Type	Number of mounts
Location	Filling status
Time of first assignment	Data volume
Last mounting date	Fragmentation
Last writing date	Closing information
Serial number	Integrity check information
External location code	

The barcode of a tape will be displayed in a tooltip if you hover the mouser pointer over the tape name under a tape pool in the **Storage** view:

Figure 11.9. Tape in a Pool – Barcode Tooltip



The maximum supported length for a barcode in Active Circle has been set to 32 characters. Any barcode characters beyond this will be ignored.



Note

The **Full** field indicates the filling status of the tape. If this box is checked, it means that no additional files should be added to the tape, although any existing files can be modified. The tape is marked full when the remaining space on the tape reaches the *filling threshold* defined for the drive type. In Active Circle, this threshold is reached when 5 GB of space remains on the tape. This is true for all tape drive types except **IBM LTO5**. For tapes in IBM LTO5 drives, the filling threshold is 15 GB.

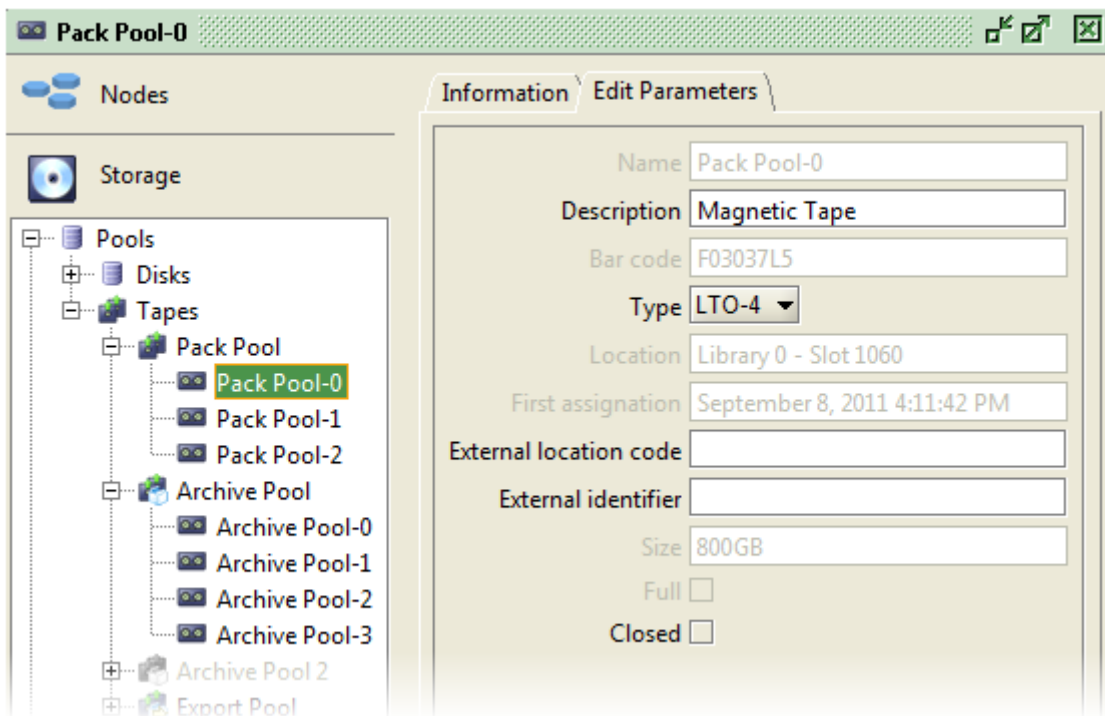
The **Volume** section shows information about the data on the tape, including the number of files, data size and fragmentation information. The **Size** field indicates the total capacity of the tape, which is automatically detected according to the specified tape type. The **Used space** field shows the space the data occupies on the tape, after compression. The **Data size** field indicates the actual size of the data on the tape, before compression. Compression is performed by the tape drive, and the compression ratio depends on the type of drive.

The **Closing Information** section shows whether the tape has been closed, when it was closed and the reason for the closing.

The **Integrity Information** section shows information relating to the last tape [integrity check](#) if the tape is in an archive pool.

Use the **Edit Parameters** tab of the tape to change the parameters:

Figure 11.10. Tape in a Pool – Edit Parameters tab



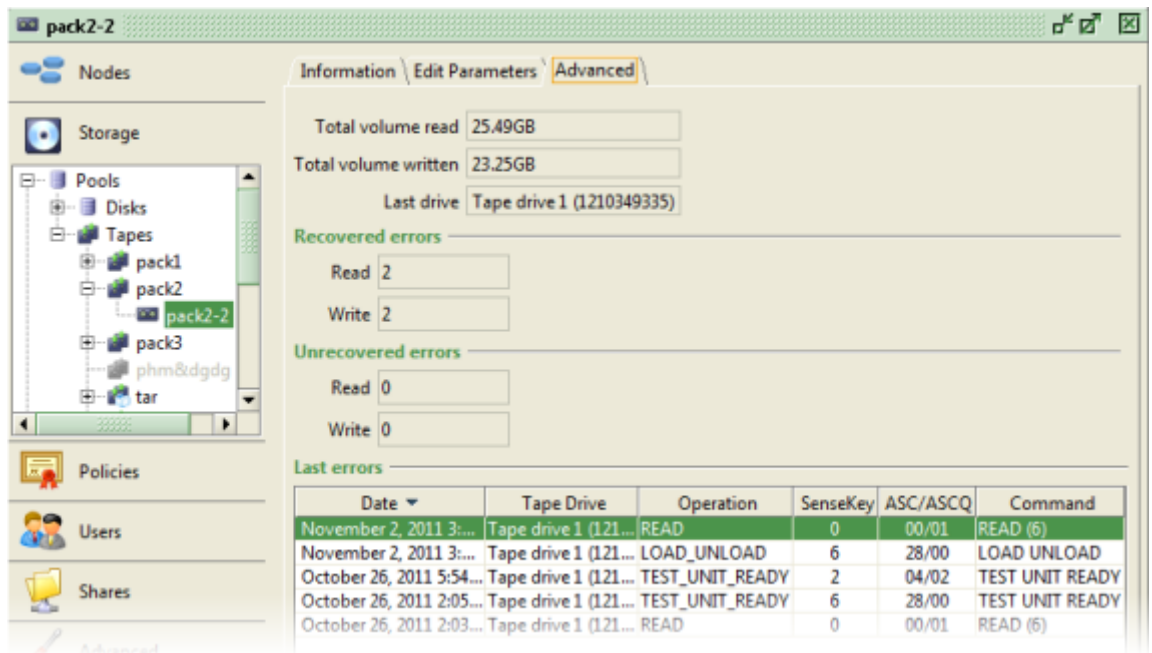
The following parameters of a tape can be modified:

Description	You can change the description of the tape here.
Type	You can change the type of the tape, depending on the compatibility of the tape drive.
External location code	A field that can be used to enter a code for helping to locate the physical location of the tape.
External identifier	Like the External location code , this field can help identify and locate the physical tape.
Closed	A manual flag that allows the administrator to indicate that a particular tape should not be modified: no additional data may be written to it and the existing files cannot be modified. The data on the tape is read-only.

Advanced

The **Advanced** tab for a tape shows information about the total volume of data that has been read and written to and from a tape since it was registered in the Active Circle system, which drive it was last used in and detailed error information for the tape. The tape drive is indicated with both name and serial number in the format `drivename (serialnumber)`.

Figure 11.11. Tape in a Pool – Advanced tab



The error sections display a summary of the number of recovered (corrected) and unrecovered errors, with separate counts for read and write errors in both categories. The list at the bottom lists details about the last errors that occurred on the selected tape.

By default, the **Last errors** list contains the last 5 errors for the tape. Each entry contains information about the date, the drive used, the operation performed (read/write/positioning), the SCSI sense data in the form of the Sense Key, Additional Sense Code/Additional Sense Code Qualifier (SenseKey, ASC/ASCQ) and the command executed. For more information about SCSI sense data codes, please see the Wikipedia article at: <http://en.wikipedia.org/wiki/ASC/ASCQ>.



Note

Two *circle properties* are related to the advanced tape information and management, and you can use these to modify the default settings for the maximum error count or age of a tape in the system.

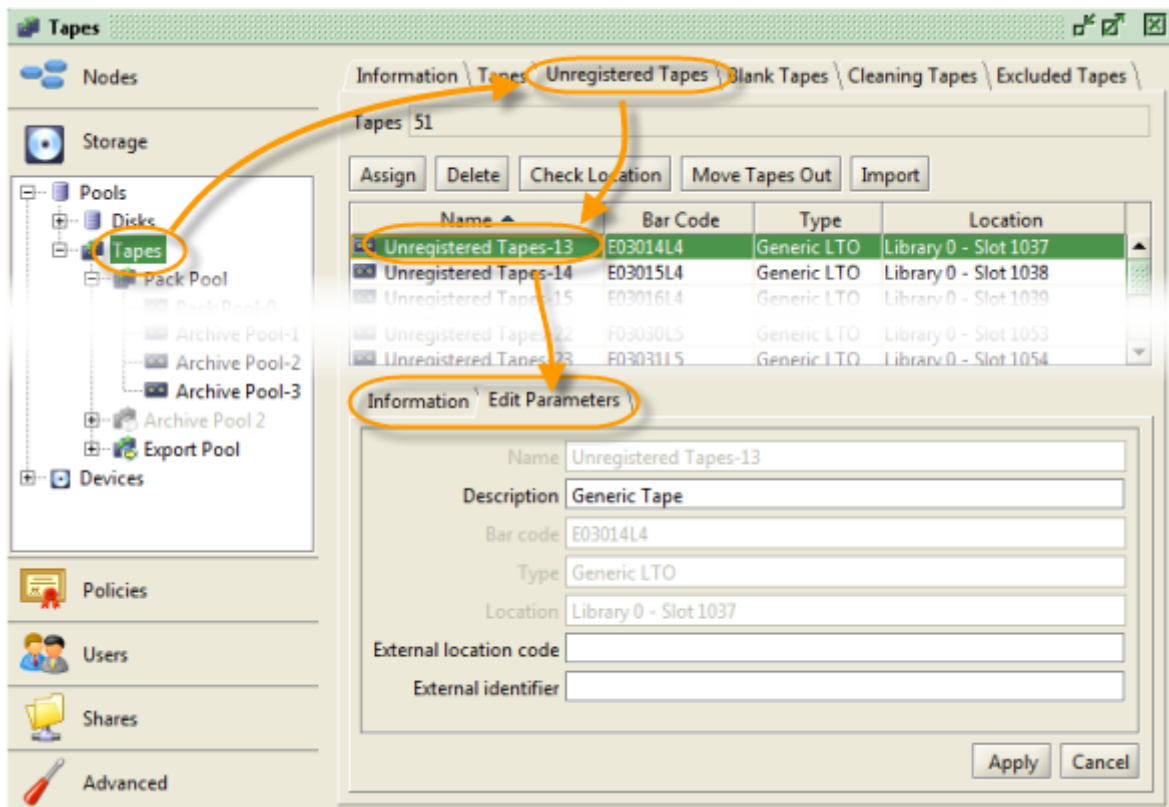
- `media.recoveredErrorsThreshold` – When the number of errors defined by this property is exceeded, the system issues a supervision note. The default value is 5.
- `media.lifeTime` – This defines the maximum time the tape can be used by the system after the first mount. If the age of the tape exceeds this value, the system issues a supervision note with a recommendation that the tape be recycled and removed from the system. The default value is 1826 days (5 years).

For information on how to modify circle properties, see “[Circle Properties](#)” in [Chapter 5, The Circle](#).

Parameters for Tapes in Default Pools

For tapes that belong to one of the four default tape pools, the tape parameters are displayed below the list of tapes when you select the tab for the default pool. For example, to see or edit the parameters for an *unregistered tape*, click on the **Tapes** item in the **Storage** view, click on the **Unregistered Tapes** tab and select a tape from the list of tapes. The **Information** and **Edit Parameters** tabs for the tape will be displayed below the list.

Figure 11.12. Edit Parameters tab for a Tape in a Default Pool



The buttons in this tab and the other tape pool tabs are explained in the section below.

Managing Tapes in Tape Pools

You can perform various operations on the tapes in the tape pools, using the options available on the **Tapes** tabs for the pools and on the **Default Tape Pools** tabs. You can assign tapes to pools, remove and delete tapes from pools, move tapes to the mailboxes, check data location, and defragment, duplicate and close tapes. These operations are described below.

Search for Tapes

You can view the list of tapes at the system level or at the pool level (including the default pools). To see all tapes, click on the **Tapes** item under **Pools** in the **Storage** view and select the **Tapes** tab. To view the tape list for a pool, select the pool in the list of tape pools in the **Storage** view and select the **Tapes** tab for the pool. You can search for tapes by applying filters that correspond to one or more of the following criteria:

- Tape name

- Tape barcode
- Tape serial number
- Tape location


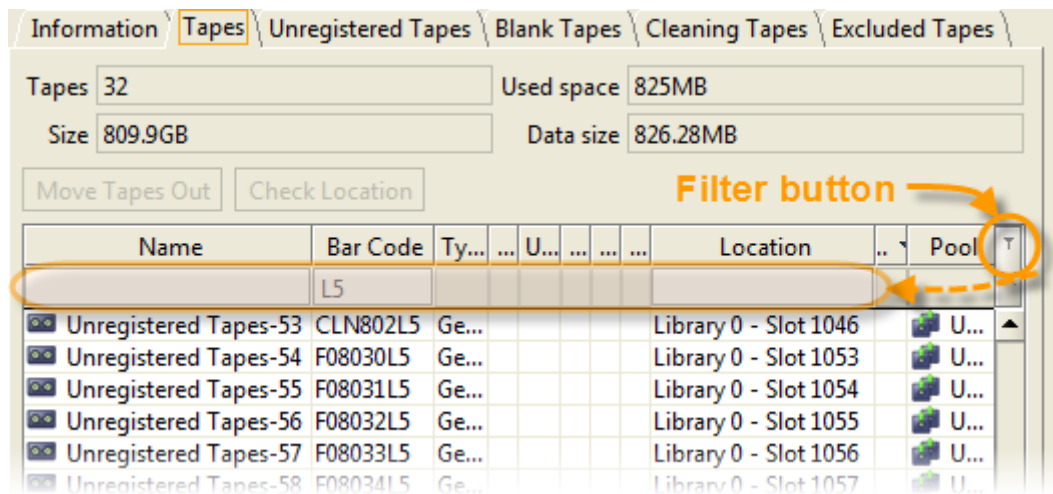
Click on the filter  button in the top right corner of the table, and type in characters in one or more of the filter fields to match any content in the corresponding columns. The list of tapes will be updated to match the criteria as you are typing. In the below example, we have typed in "L5" in the bar code filter field to find all bar codes containing those characters. You can enter characters in several columns at the same time to narrow your search further.

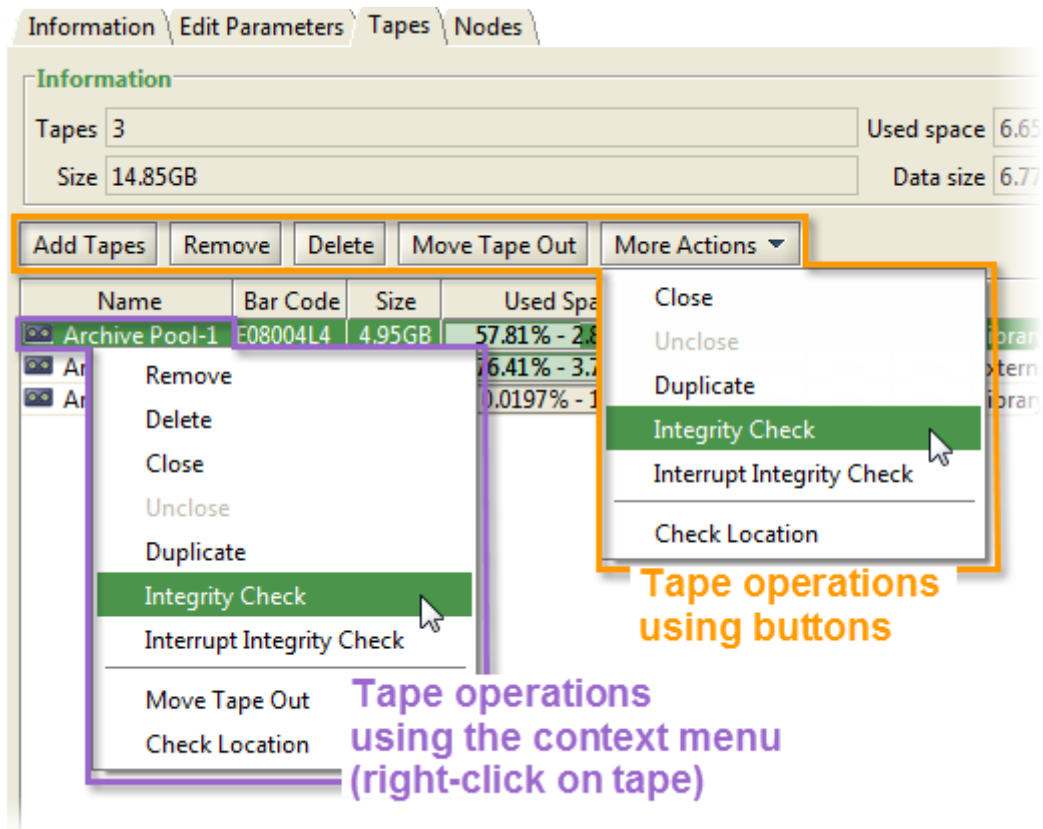
Figure 11.13. Filter the List of Tapes



Tape menus and buttons

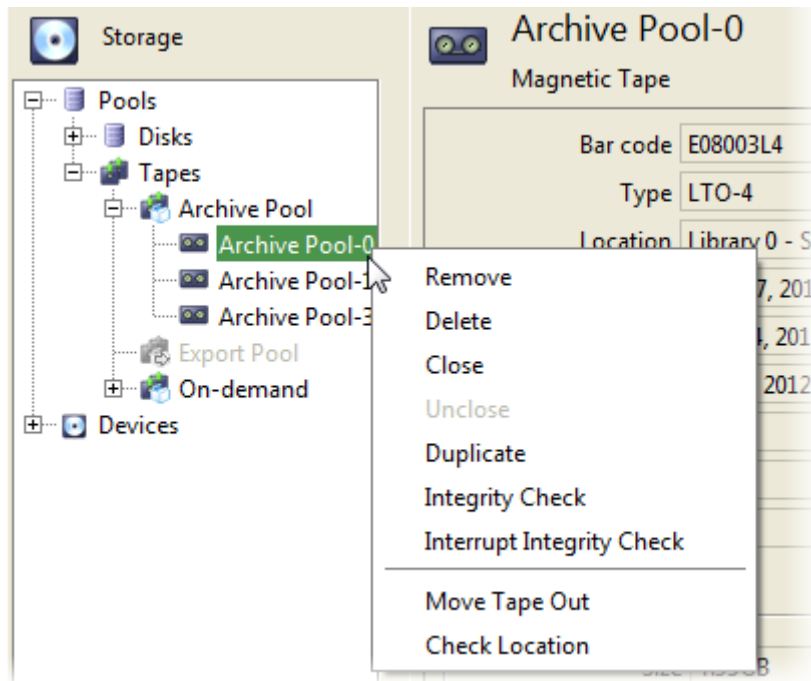
For most tape operations, you can access the controls in two ways; either by selecting a tape and using the row of buttons above the tape list, or by right-clicking on a tape and selecting an option from the context menu for the tape.

Figure 11.14. Tape Operation Choices in Archive Pools



You can also access these controls by right-clicking on a tape in a pool in the tree view:

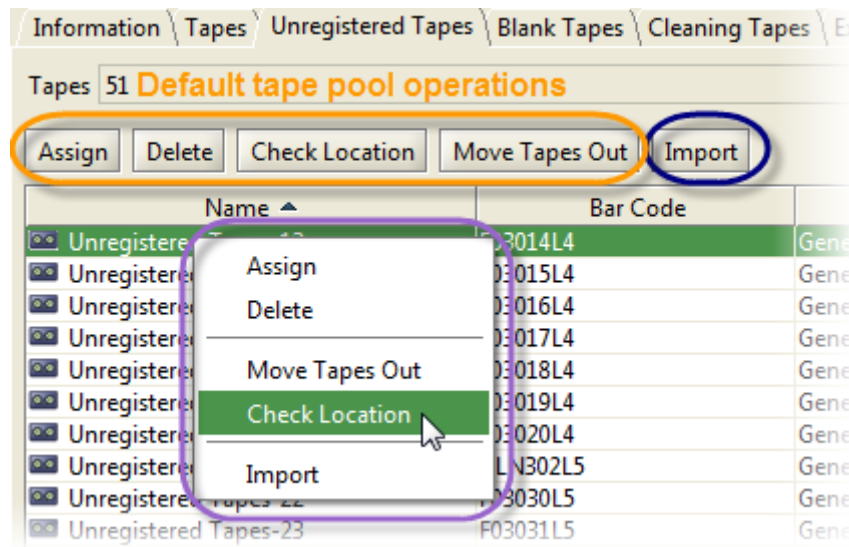
Figure 11.15. Tape Operation Choices on Tape Context Menu



For pack pools, the **Duplicate** item will be replaced by **Defragment**.

The common set of operations are also available from each of the default tape pools:

Figure 11.16. Tape Operation Choices in Default Pools



Assigning Tapes to Pools

Before your tape pools can be used to store data, you must assign tapes to the pools. Initially, all the tapes in the libraries are assigned to the **Unregistered Tapes** pool by the system. You need to re-assign them to the appropriate pools, depending on the intended use.




Tip

For tape pools with **Auto Expand** activated, the administrator only needs to assign tapes to the **Blank Tapes** pool. The system will automatically take care of assigning the tapes to the proper pools when necessary. For auto-expand pools, it is not necessary to assign tapes directly to each pool.

To assign a tape or a group of tapes to a tape pool from one of the default pools, select it (you can select more than one) and click the **Assign** button. The **Assign Tape** selection dialog box is displayed.

Figure 11.17. Assign Tape




The default pool for assigning tapes is **Blank Tapes**. Select the destination tape pool for the tape(s) and the **Tape type**. When you have finished, click **OK** to close the dialog box

and move the tape(s) to the appropriate tape pool. The tapes are removed from the list of **Unregistered Tapes** and should now appear as new tapes in the designated tape pool.

You can also assign a tape to a pool by *adding* it to an already selected pool, without having to first go to the **Unregistered Tapes** tab. Select the pool you want to add the tape(s) to, and click on the **Add Tapes** button on the **Tapes** tab. Alternatively, you can right-click on the pool in the **Storage** view and select **Add Tapes** from the menu. The **Add Tapes** selection dialog box is displayed.

Figure 11.18. Add Tapes



Select the tape pool you want to move the tapes from (**Unregistered Tapes** is the default), and select one or more tapes to add from the list. Click on **Assign** and the **Assign Tape** dialog is displayed where you can select the tape type. Click on **OK** in that dialog to add the tape(s).



Important

Make sure that you assign the appropriate cartridge to the appropriate pool. You should not assign a cleaning tape to a pool used for archiving or a normal tape to the pool of **Cleaning Tapes**. By default, if the barcode starts with the letters "CLN", it will be considered as a cleaning tape by Active Circle. You can change or add barcode prefixes to identify cleaning tapes by modifying the `media.excludedTapeBarcodes` circle property. For more information on editing circle properties, see ["Circle Properties"](#).

Although there is no limit to the theoretical capacity of a tape pool, it is recommended that you only add tapes to pools as they are needed; for example in groups of 10 or 20.

Deleting or Removing Tapes

All **Tapes** tabs for tape pools contain one or both of the options **Remove** and **Delete** that operate on the tape selected in the table. These options are used to move tapes back to either **Unregistered Tapes** or **Blank Tapes**.

- Remove** Dissociates the tape from the tape pool and places the tape in the **Blank Tapes** pool.
- Delete** Deletes the tape from the tape pool and returns it to the list of **Unregistered Tapes**.



Warning

Using either the **Remove** or **Delete** option will delete all data stored on the tape.



Note

When you delete or remove a tape from a pool, all the metadata associated with the tape are cleaned on all the nodes. This process can take some time. During this time the tape still appears in the pool but its name will be displayed in italics. It will remain in this state until everything is cleaned on all the nodes. This means that if one of the nodes is not running, the tape will remain in this state until the node is started and the removal process can be completed (synchronized) on that particular node.

If you have a configuration with two nodes on the same server (two Active Circle services running on the same system, like in the Active Media Library), the tape removal will not be completed until both services are running.

Checking Location Information

Using the **Check Location** button, you can refresh the location information for tapes in the library. In pack or archive pools, the **Check Location** option is found under the **More Actions** button.

Moving Tapes to Mailboxes

If you need to remove a tape from a library, the tape must first be moved to one of the library mailboxes. Use the **Move Tapes Out** button to move the selected tape(s) out of the library and place them in the mailboxes (physically moves the tapes from their currently assigned slots in the library to the mailbox slots). The tapes will remain assigned to their pools, but they will have the status *external*. You cannot move more tapes than there are available mailbox slots. For information on moving tapes from mailboxes back into libraries, see [“Moving Tapes To and From Mailboxes”](#).

Closing a tape

The administrator can decide that a particular tape should have no more data written to it. This is done by selecting a tape, clicking on the **More Actions** button and selecting **Close** from the [menu](#). The tape becomes read-only, and cannot be modified. The tape can be opened for writing again by selecting **Unclose** from the same menu.

Integrity Check

You can perform an integrity check of tapes used in archive pools. This is based on a signature that is calculated for each TAR file created during an archiving operation.

Chapter 11. Tape Pools

The signature is stored in the metadata associated with each TAR file. The verification procedure recalculates the signatures while reading the files and compares them with the initial signatures from the metadata.

To perform an integrity check, select one or more tapes in an archive pool, right-click to bring up the [context menu](#), and select **Integrity Check**. A message will be displayed asking you to confirm the operation (if you select many tapes, the process can take a while). Click on **OK** if you want to continue. During processing you can abort the checking by selecting **Interrupt Integrity Check** from the tape context menu.



Note

No integrity check is available for tapes in LTFS tape pools. The functionality only applies to tapes with TAR format.

If any invalid checksums for TAR files are detected, a supervision note will be issued for each tape with errors, giving details about the corrupt TAR file(s) and the associated pool(s). The integrity status of the tape will be set to "FAILED" and the processing of the tape will end.

By default, 10% of each tape will be verified. You can change this percentage by modifying the value of the *circle property* `archive.integrityCheckPercent`. This defines the volume of TAR files (in percent) to check during tape integrity check. The value must be between 5 and 100. The change will be applied after restart of the node. For more information on how to set values for circle properties, see ["Circle Properties"](#).

You can verify the results of the integrity check in two ways:

- For one tape: Click on a tape under an archive pool in the **Storage** tree view and select the **Information** tab for the tape. At the bottom of the information pane, the date and result of the last integrity check is displayed:

Integrity Information

Last integrity check date	May 31, 2012 11:02:46 AM
Last integrity result	VALID

- For several tapes: Click on an archive pool or the parent **Tapes** item in the **Storage** tree view and select the **Tapes** tab. This will display the list of tapes at either the pool or the system level, depending on the item you just selected. By default, the table on this tab does not display integrity check information for the tapes. To add this information to the table, right-click on a column header and select [Column Properties](#). In the **Column Display Properties** dialog, place a check mark next to the column labels **Last integrity result** and **Last integrity check date**. Click on **Close**, and the selected columns will be displayed for each tape in the list (scroll to the right if the columns are not visible):

Name	Bar Code	Type	Last integrity result	Last integrity check date
On-demand-0	E08007L4	LTO-4	VALID	May 11, 2012 4:36:54 PM
Archive Pool-1	E08004L4	LTO-4	VALID	May 31, 2012 11:02:27 AM
Archive Pool-0	E08003L4	LTO-4	VALID	May 31, 2012 11:02:39 AM
Archive Pool-3	E08001L4	LTO-4	VALID	May 31, 2012 11:02:46 AM

More detailed information on the integrity check can be found in the log of the node that performed the check.



Tip

You can also perform tape integrity checking using the Active Circle command line interface with the command **accheck --tape**. This can be useful if you want to automate tape integrity verification and reporting by using scripts.

Unused Tapes Check

Verification of unused tapes is an automatic tape operation performed by the system at regular intervals. Tapes which have not been used for a certain period (6 months by default) will be loaded in a drive to check if they are still usable. Offline tapes are by default not checked.

Using [circle properties](#), you can change the default behavior of this feature:

Tape Action	The default action is to load ("Mount") the tape in a drive. You can disable the automatic checking by setting the tape action to "None". To perform an integrity check of the tape, set the value to "Integrity". <i>Circle property:</i> <code>media.unusedTapeAction</code> <i>Default value:</i> Mount
Idle Period	You can define the length of the period of inactivity for a tape before it is checked. <i>Circle property:</i> <code>media.unusedTapeCheckPeriod</code> <i>Default value:</i> 6 months
Offline Tapes	You can define that offline tapes should also be part of the automatic checking for unused tapes, by setting the value for the following property to "true": <i>Circle property:</i> <code>media.unusedOfflineTapeCheck</code> <i>Default value:</i> false



Note

This functionality only applies to tapes which belong to a real Active Circle tape pool and which contain data. Unregistered and blank tapes will not be checked.

Defragmenting and Duplicating Tapes

Active Circle implements a feature for defragmenting or duplicating the tapes used in tape pools. The defragmentation functionality applies to tapes used in *Pack* pools, while the duplication function is used for tapes in *TAR Archive* pools.



Note

Tapes in Export pools cannot be duplicated.

Defragmentation and duplication works by copying the data directly from one tape to another. Because of this, these operations require two tape drives in order to function.

To access defragmentation, right-click on a tape in a *pack* pool and select **Defragment**. Alternatively, you can select the tape, click on the **More Actions** button and select **Defragment** from the menu.

To access duplication, right-click on a tape in a *TAR archive* pool and select **Duplicate**. Alternatively, you can select the tape, click on the **More Actions** button and select **Duplicate** from the [menu](#).



Important

After defragmentation or duplication is completed, all the data on the source tape (the one being copied from) will be deleted.

For tapes in pools, an indicator of the level of fragmentation for a given tape appears in the **Information** tab. You can set a [fragmentation threshold](#) for each tape pool.

Importing an Archive

The tape import functionality allows an administrator to create a file structure on an Active Circle share (preserving the paths, filenames, UIDs, GIDs, dates, etc.) that corresponds to the contents on tapes external to the Active Circle storage system. The files and folders will be visible on the share as if they had been deposited there by a user in a regular way. This is useful for importing data on tape from outside the Circle. For TAR archives, a new folder will be created for each imported TAR file. For LTFS archives, the current file structure will be imported.



Note

The import feature can also be used to import data from a tape created by Active Circle (the system uses the tape label to identify the origin). In this case, each file created on the share will have the name of the original share as the top level directory in its path.

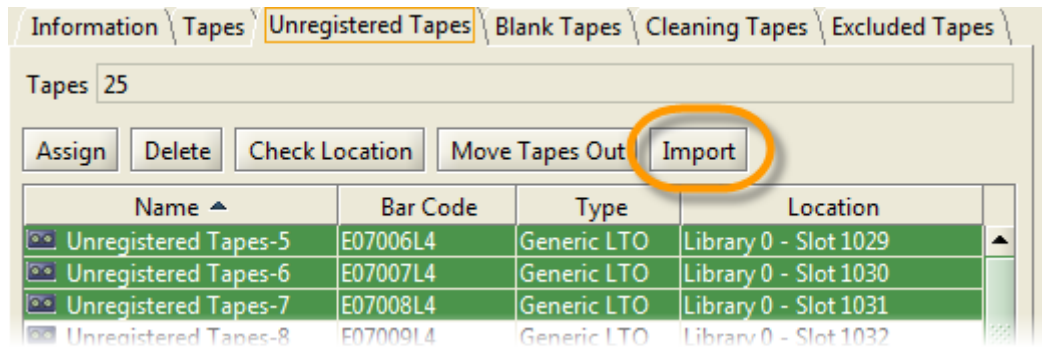
This functionality is only available if the archiving operations have not been disabled for the Circle (the *circle property* called `archive.disable` must be set to *false*, which is the default setting). The user performing the operation must have one of the profiles of **Data administrator** or **Operator**.

Importing a tape archive requires the existence of at least one share meeting the following conditions:

- The share must have an Archiving Policy corresponding to the tape type (TAR or LTFS)
- The Archiving Policy must have at least one rule that is not an on-demand archiving rule
- You must have write permission to the share

Before the import is started, the tapes to import from must be loaded into the tape library. Then the library must be synchronized using the Active Circle Administration Tool, by selecting the library in the **Storage** view and going to the **Inventory** tab and clicking on the **Synchronize** button. The tapes will then be available in the **Unregistered Tapes** list.

Figure 11.19. Import Button in Pool of Unregistered Tapes



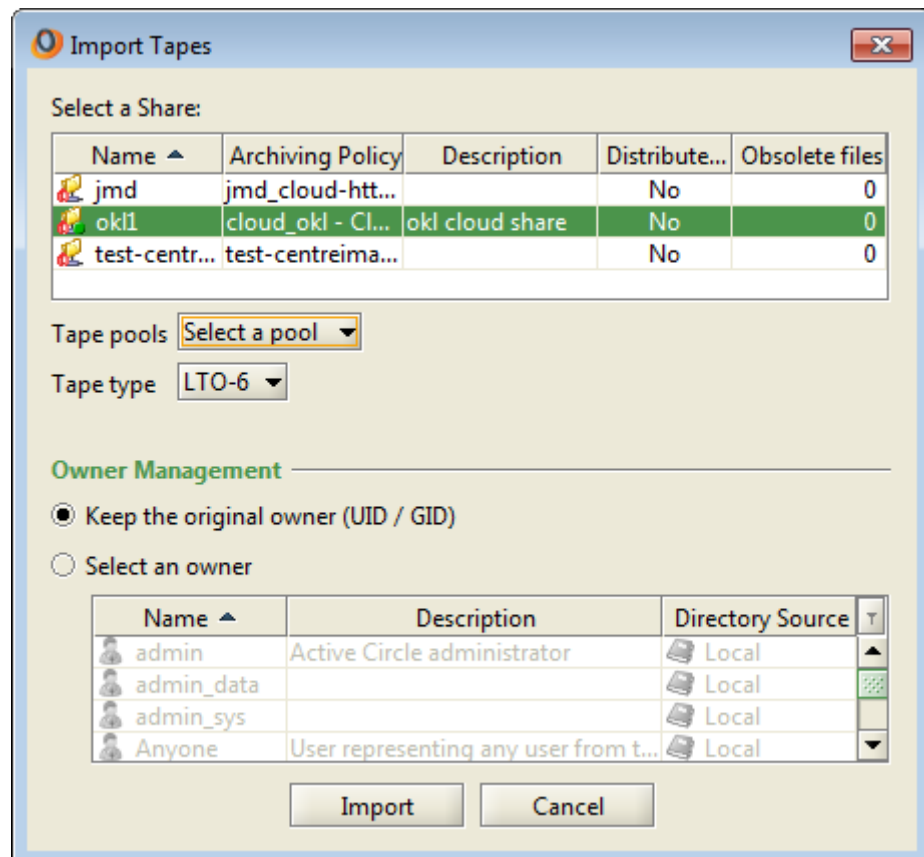
To start the import process in the Administration Tool, click on the **Tapes** item in the **Storage** view. Select the **Unregistered Tapes** tab, select the tape(s) to be imported from the list of tapes, and click on the **Import** button. The **Import Tapes** dialog box is displayed.



Important

All the tapes selected in an import operation must have the same generic tape type.

Figure 11.20. Import Tapes dialog



Select the **Share** where the files should be made available after import. Only shares that fulfill the requirements stated above will be available.

Select the destination **Tape pool** for the imported tape(s). Only pools defined with a standard constraint in the archiving policy associated with the selected share will be available.

Select the **Type** of the tape(s) to be imported. Only types compatible with the tapes' generic type will be available. A drive supporting this type must be present in the library. If the type selected is incompatible with the available drives, a supervision note will be issued.

Specify whether to use the original UID/GID information from the tape files (the default choice) or click **Select an owner** and select a user who should be the new owner of the files. Click on OK to start the import process. When the import is finished, the tape should appear on the list of tapes for the specified pool. The imported files should be available on the share selected during import, with the same structure as in the original data on the tapes.

Export Pool

An *Export Pool* is used to create an archive on tape for the purpose of removing the archived files from the library. This can be a useful feature for making external backups and moving data between different circles or systems.

When creating an export pool, you can choose between two formats for the data on the tapes:

- **Posix TAR**: The default archiving format of the Active Circle system, based on the Posix standard.
- **LTFS**: The archive will be created using the Linear Tape File System, to be read on systems supporting the LTFS format.

You may not change the format of the pool once it has been created. To create an export pool, see [“Creating and Activating a Tape Pool”](#).

The actual export operation is launched using the Active Circle **File Explorer** application or the **acarchive --export** command on the Active Circle Command Line interface. There, the user selects the files to export and starts the archiving operation (please refer to the *Active Circle File Explorer Guide* or the *Active Circle Command Line Guide* for more information on these procedures). However, for the export to be possible, the administrator must ensure that the prerequisites for this functionality are satisfied, which must be done using the Active Circle Administration Tool. At the end of an archiving for export, the administrator or an operator should remove the tape(s) with the exported files from the tape library. The requirements to prepare for this are described below.

Prerequisites:

- The *circle property* `archive.disable` must be set to *false* (default)
- The *circle property* `archive.enableExport` must be set to *true* (this is set to *false* by default). For more information on modifying this, see [“Circle Properties”](#).
- The user performing the export must have the **Archive** (A) access right on the share to export from.
- The administrator must create an archive *Export Pool*, activate it on a node and assign tapes to it, as described earlier in this chapter.
- The archiving policy that is associated with the share must have the **Export Pool** selected in the Advanced Policy Parameters (see [“Advanced”](#)).
- Only tapes conforming to the LTO-5 specification or higher can be used for export in LTFS format.

When the above steps are completed, users can perform archiving for export using the File Explorer.

Chapter 11. Tape Pools

The export pool will appear under the **Tapes** item under **Pools** in the **Storage** view, along with all the other tape pools. The parameters and operations available in the tabs for an export pool are the same as for other archive tape pools. However, a regular TAR archive pool may not be used for exporting archives using the File Explorer.

Chapter 12. Cloud Pools

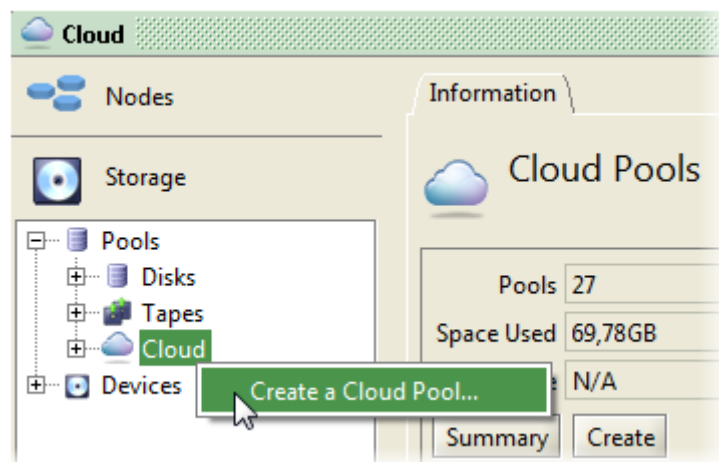
A cloud pool is storage space located off-site in a hybrid cloud storage solution. The locally stored data on Active Circle shares can be added to the cloud storage through association with an archiving policy that uses a cloud pool as data location.

In Active Circle, the cloud pool refers to a location already defined in the cloud storage of the external provider. The cloud storage will be set up by the provider and the connection information provided to the Active Circle administrator who is creating and managing cloud pools in the local part of the Active Circle Storage System. Cloud pools do not consist of local physical data partitions like disk or tape pools, but have pointers to external storage in the cloud. The archiving policies and shares associated with the cloud pools define which data will belong (or be transferred) to the cloud storage. This chapter explains how to create a cloud pool and manage it.

Creating a Cloud Pool

To create a new cloud pool, right-click on the **Cloud** element in the **Storage** view. Select **Create a Cloud Pool** from the context menu. Alternatively, you can use the **Create** button on the **Information** tab for the **Cloud** element.

Figure 12.1. Create Cloud Pool menu item



Note

Before you can proceed with creating the cloud pool, you will need the address of the cloud server and the login information which is provided by the cloud storage provider.

The pool creation dialog will be opened.

Figure 12.2. Cloud Pool Creation dialog

Cloud Pool Creation

Name: Cloud Pool 1

Cloud Settings

Provider: oodrive

Server URL: https://circletest1

Login Name: testuser

Password:

Confirm password:

Activation Nodes :

Name	Type	Description
JMD450	Domain	
circletest1.active-circle.Node	active-circle.Node	

Ok Cancel

You must enter a name to identify the new cloud pool, specify connection information and add nodes:

- | | |
|--------------------------------|--|
| Name | The name cannot include special characters or spaces. Numbers, underscores, hyphens and periods are allowed. |
| Provider | This drop-down list lets you select the provider who is hosting your cloud storage. By default this is predefined the default hosting partner of Active Circle (oodrive), but if your cloud agreement involves another provider, it may be selectable in this list. |
| Server URL | Enter the URL address for the server where your cloud storage is hosted. This information will be supplied by your provider. |
| Login name and password | Enter your username and password for the cloud server. This information will be supplied by your provider. If you do not specify a password, the program will ask if you want to proceed without a password. If you accept this, the pool will still be created in the user interface, but you will not be able to activate it on any nodes, and no data partition will be created in the cloud. In this case you may enter the password later on the Edit Parameters tab of the pool. |
| Activation Nodes | A cloud pool cannot be used unless it is activated on one or more nodes in the local part of the hybrid cloud solution. You can select the activation nodes in this dialog, or you can do it later as described further |

down in this chapter. Use the **Ctrl** key to select more than one node from the list. Note that until the pool is activated on nodes, it's entry in the list of pools in the **Storage** view will be greyed-out.

During pool creation, a logical data partition is created locally. However, the actual data partition will be created in the cloud only upon the initial activation of the pool on an Active Circle node. If activation fails, the pool will be taken off-line and the availability will be indicated as "Error". There will be one or more supervision notes issued in this case, so please consult the notes as described in [Chapter 32, Supervision](#). In addition, you may look in the [node logs](#).



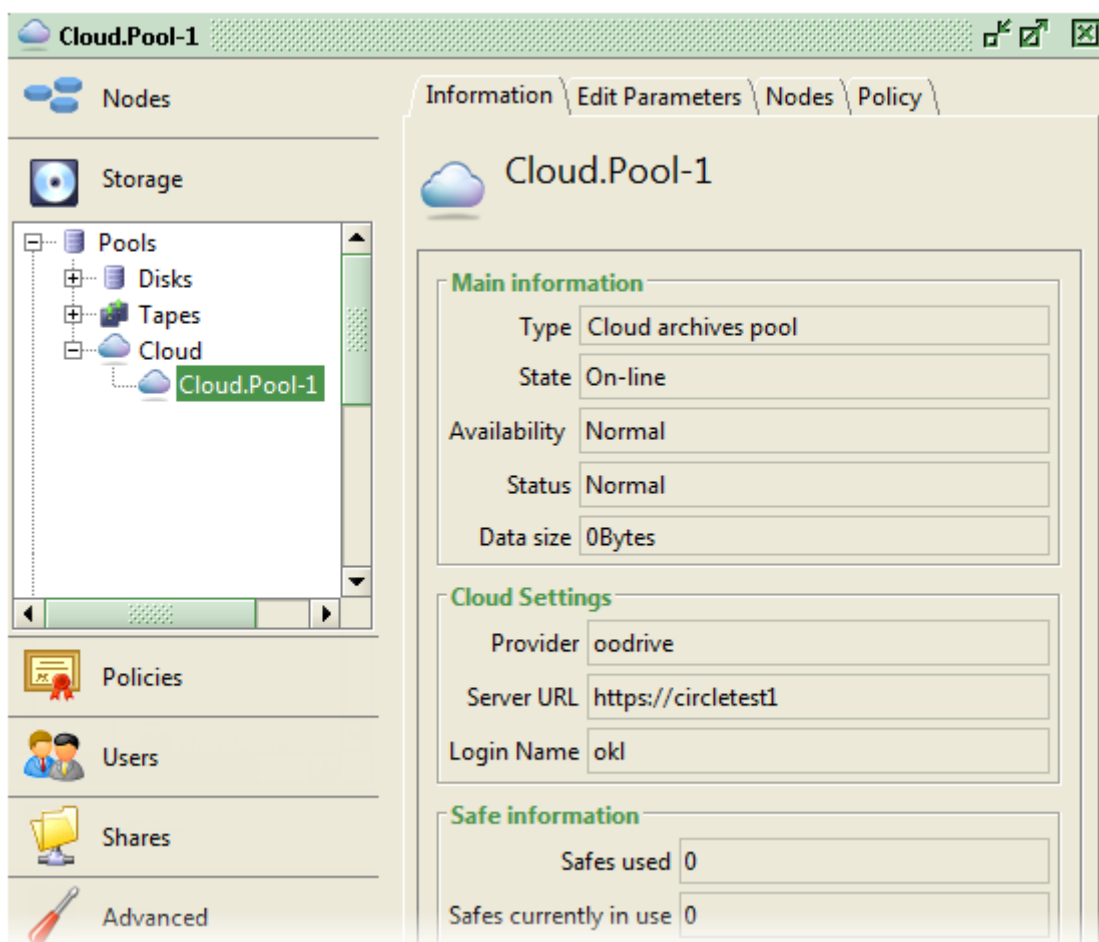
Note

Before you can archive data to the cloud, the cloud pool must be associated with an archiving policy that is active on a share.

Cloud Pool Information

The **Information** tab displays all the information about the selected pool. Most of the information is similar to other pools, but cloud pools do not have information about physical storage like data partitions or tapes.

Figure 12.3. Cloud Pool Information tab



Main information

Type	Indicated that this is a cloud archives pool.
State	Current state of the disk pool. Possible values are: On-line, Off-line.
Availability	Current availability of the cloud pool. Possible values are: Normal, Read only.
Status	Current status of the disk pool.
Data size	Size of the data which have been archived in the cloud pool.

Cloud Settings

Provider	The provider hosting the cloud storage.
Server URL	Server address for the cloud storage.
Login Name	User name on cloud server.

Safe information

Safes used	Number of safes to which data has been written.
Safes currently in use	Number of safes currently open for use.

Cloud Pool Password

To enter a new password, go to the **Edit Parameters** tab for the pool and click on the **Modify Password** to the right of the **Login Name** text field. Type in the new password twice in the password dialog box and click on **OK**.

Figure 12.4. Setting Cloud Pool Password

The screenshot shows a web-based configuration interface for a cloud pool. At the top, there are four tabs: 'Information', 'Edit Parameters', 'Nodes', and 'Policy'. The 'Edit Parameters' tab is currently selected. Below the tabs, there are several configuration fields: 'Name' (set to 'Cloud.Pool-1'), 'Description' (set to 'My Cloud Storage'), 'State' (a dropdown menu set to 'On-line'), 'Availability' (a dropdown menu set to 'Normal'), and 'Login Name' (set to 'testuser'). To the right of the 'Login Name' field is a button labeled 'Modify Password', which is circled in orange. The background of the form is a light beige color.

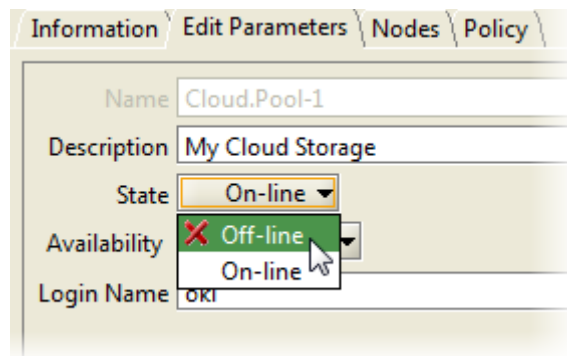
The password will be supplied by the cloud provider. When you receive a new password, you enter it as described above.

Cloud Pool States

It is possible to modify the state of a pool by taking it off-line in order to prevent any read or write activity on the pool, something you might do during a maintenance operation on the server.

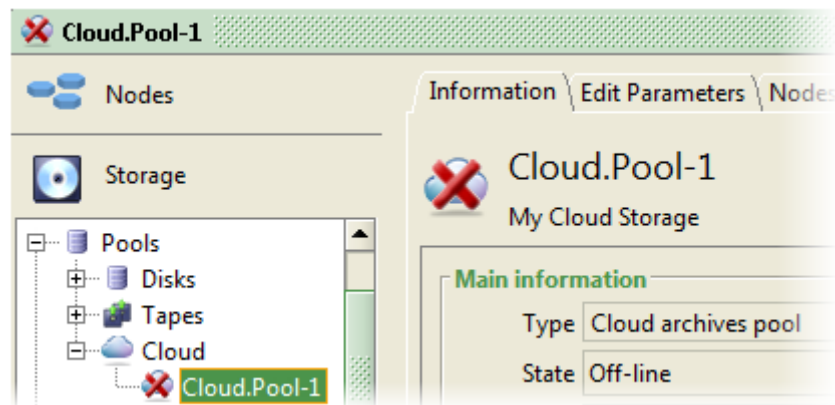
To take a pool off-line, in the **Edit Parameters** tab for the pool, click on the **State** drop-down list, select **Off-line**, and click **Apply**.

Figure 12.5. Set Cloud Pool Off-line



A red 'X' appears on the pool in the tree view, the window heading and the information pane to indicate that it is off-line:

Figure 12.6. Off-line Pool



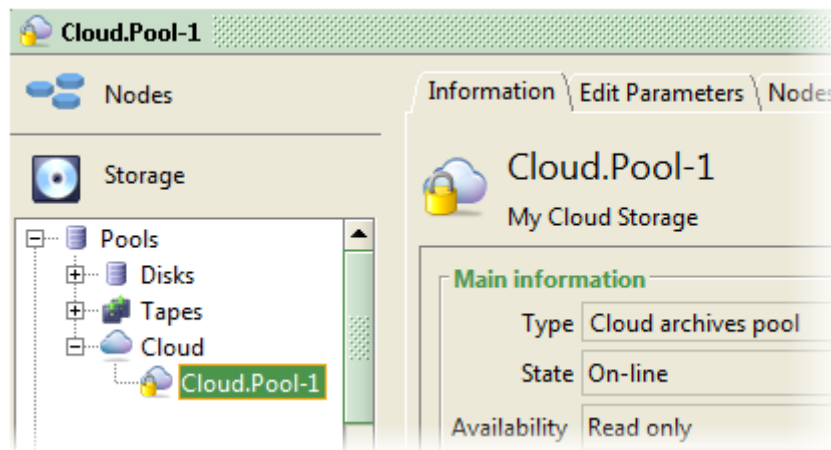
To return the pool to the on-line state, perform the same process selecting **On-line** and clicking **Apply**.

Cloud Pool Availability

You can modify the availability of a pool by making it read-only.

To make a pool read-only, in the **Edit Parameters** tab for the pool, click on the **Availability** drop-down list, select **Read only**, and click **Apply**. This change is reflected in the **Availability** field that appears in the **Information** tab and in the yellow lock that appears on the icon of the pool in the administration interface.






Figure 12.7. Read-only Cloud Pool



Cloud Pool Status Indicators

If an error occurs with the cloud pool, the icon for the pool will add a red dot with a white exclamation mark in the lower left corner. If the pool is off-line as well, the icon will be crossed out in red in addition to showing the red error dot. This can be combined with a black wrench indication that the pool is in maintenance. Several indicators can be combined to display more than one pool condition at the same time, as shown in the table below.

Table 12.1. Cloud Pool Status Indicators

Indicator	Meaning
	Cloud Pool in normal operation, no problems detected.
	The pool is <i>Off-line</i> .
	The pool is <i>Read-Only</i> and <i>Off-line</i> .
	There is an error with the pool, and the pool is off-line. A <i>Supervision Note</i> will be issued for the error, to help you troubleshoot the problem.
	The pool has been placed under maintenance. This could be because the nodes the pool is activated on, are in maintenance mode.

Node and Policy Information

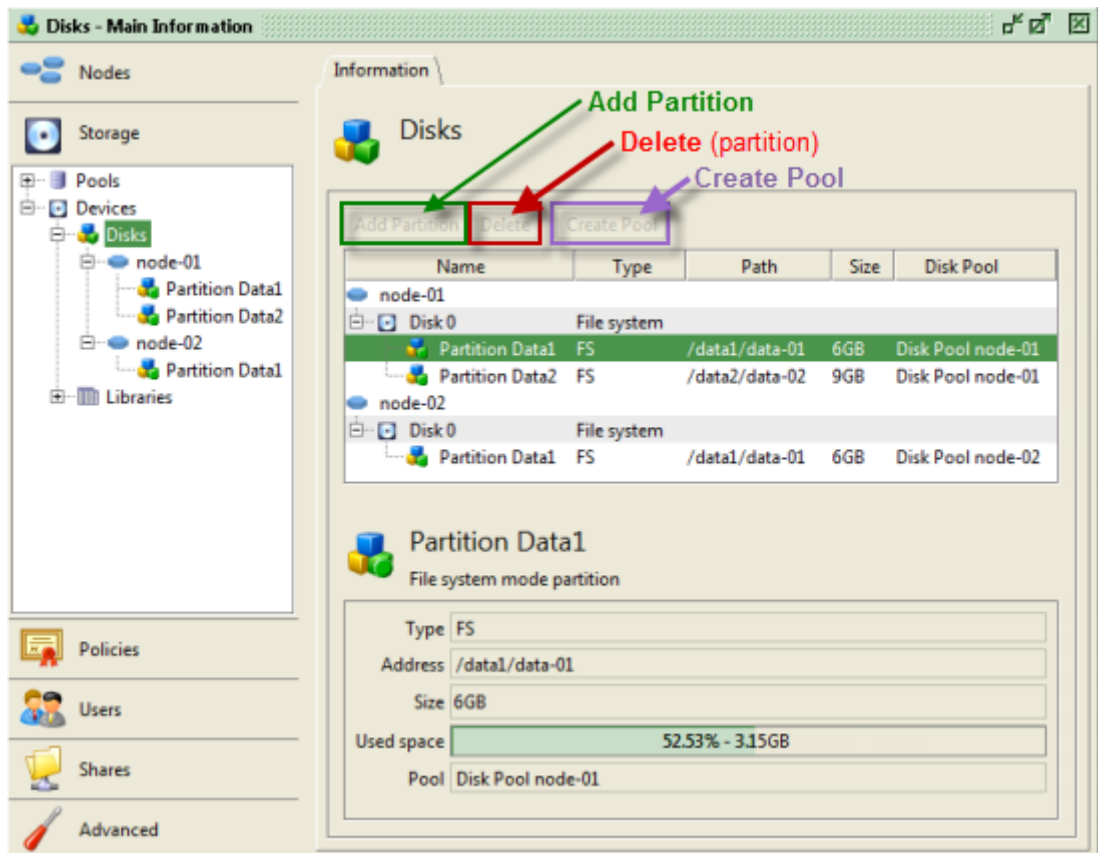
Each cloud pool has two tabs showing node and policy information, just like for any other storage pool type. The **Nodes** tab shows the node(s) the cloud pool has been activated on. On this tab, you can deactivate or activate nodes for the pool. The **Policy** tab lists the policies which are using the pool.

Chapter 13. Disk Devices

The *Disks* are the storage devices containing the *Partitions* used in the Active Circle *Disk Pools*. The partitions, which need to be predefined on disk storage, are the storage units that make up the disk pools the system uses for storing data.

When the Active Circle system starts up, it detects any hard disks attached to the system and lists them in the table located under **Storage** view > **Devices** > **Disks**. Here you can register the storage partitions to be used by the system.

Figure 13.1. Disks – Information tab



You can display the same information for one node at a time, by selecting a node name under the **Disks** item. Selecting a disk or partition from the list will cause information about the object to be displayed under the list.

Initially, only a list of disks is displayed, no partitions are visible. Even though the administrator has defined storage partitions for use with Active Circle on the disk storage attached to the Circle, they must be explicitly added in the Administration Tool to be visible and usable by the system. These procedures are described in the following sections.

Partition Types

There are two main types of partitions, defined by how they are accessed by the Active Circle Nodes:

Local Partition

File system partitions on disk storage which is attached directly to a specific Active Circle node server. A local partition needs to be included in a local disk pool, which will be available only to the node the storage device is attached to. A local partition is defined and mounted on the local server by the operating system when the server starts. A special type of local partitions are *raw partitions*, which will be described below.

Shared Partition

Shared partitions are file system partitions located on disk volumes in a SAN architecture, or NFS volumes exported on the network. Shared partitions need to be included in shared disk pools for use in Active Circle, and they will be available to all the nodes that belong to the shared pools. Shared partitions are defined in a similar way to local partitions, but they are *not* mounted by the operating system when the server starts. Mounting of shared partitions is handled automatically by the Active Circle software on the server that the system selects as the *leader node* in the shared disk pool.



Important

For partitions to be available for addition to the list of storage units available in the Administration Tool, and thereby for inclusion in storage pools, the partitions must first be configured on the storage hardware attached to the Circle. This is normally done before installing the Active Circle software. The guidelines for creating, formatting and configuring local and shared partitions are therefore described in the prerequisites part of the *Active Circle Installation Guide*. Please refer to this manual if you have not yet configured partitions, or you wish to add new storage hardware to the system.

Adding a Local Partition

You should define a partition in the Administration Tool for each of the data partitions created for use with Active Circle on your storage hardware.

There are two possibilities for writing data to local partitions:

File System (FS) mode

The file data is managed by the operating system's file system. This is the normal mode of operation. Active Circle recommends using the **ext4** file system for Linux servers.

Raw mode

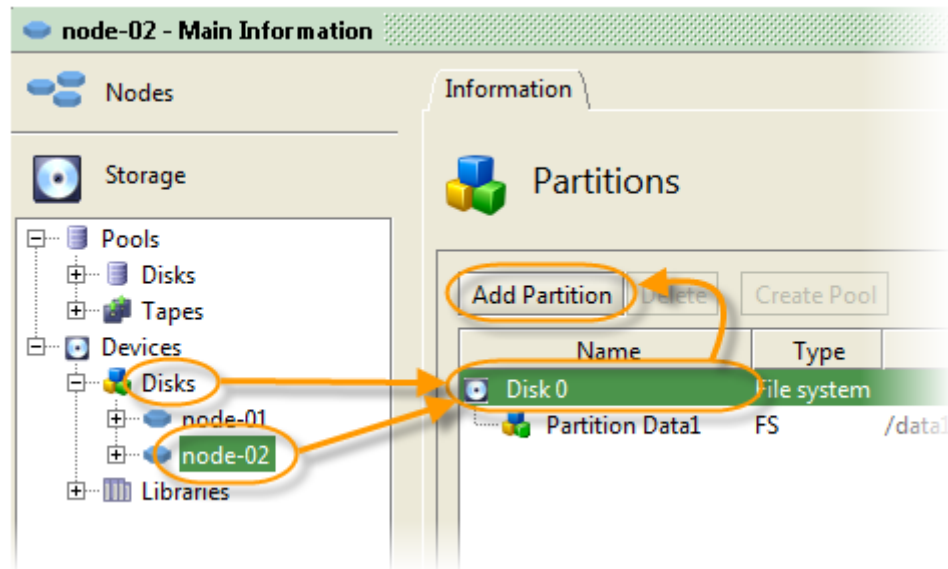
The file data is managed by Active Circle, which controls the peripherals without the use of a file system.

Add a Partition in File System mode

A new partition must be declared in an *empty directory* of a preconfigured file system partition. This means you may have to create a new, empty directory under the mount point of the partition you want to add, as the partition may have some data on the root level placed there by the operating system. For example, if the partition you want to add is mounted as `/data1`, you should create a subdirectory `/data1/data` to use as path for the partition.


To add a partition, select either the **Disks** item or the node name the disk is located on, in the **Storage** view. For each node, all local file systems will be grouped under one common item, usually "Disk 0". Select the disk item and click on **Add Partition**.

Figure 13.2. Add a Partition button



The **Add Fs Partitions** dialog box is displayed.

Figure 13.3. Add a Partition dialog



Specify a partition name, a path to an empty directory, and the size of the partition in Gigabytes. It is recommended to specify the size of the entire partition. The simplest way to do this, is to check the box **Use free space**, to let the system determine how much space is available on the volume and use all of it for the partition. The minimum allowed size is 10 GB. If only 10 GB of free space remains, the system will consider the partition as being full. Active Circle does not recommend dividing a file system partition into multiple data partitions, or letting other applications use portions of the same partition. Click **OK** to create the partition. The partition appears in the table below the disk with which it is associated and under the node it belongs to in the tree structure in the **Storage** view.

**Tip**

To determine the correct size of a file system partition, issue the command **df -m**, which will give you more precise partition sizes in megabytes (graphical disk tools or **df -h** may report more available space than there really is). Take the number in the **Available** column for your partition, convert it into gigabytes (divide by 1024) and deduct 1 gigabyte as buffer. For example, if the command returns a value that you convert to 151 GB, specify 150 GB as the size of the partition.

Add a Partition in Raw mode

To activate the Raw mode, you will have to edit the file `/activecircle/.localvars` for the Node where the partition is to be located, adding (or uncommenting if they already exist) the following lines:

```
SETENV CELL_PROBE_SCSI          true
SETENV CELL_PROBE_SCSI_DISK_DRIVES  true
```

In order for these changes to take effect, you must restart the Active Circle service on the node (see [Chapter 3, The Active Circle Service](#)).

**Important**

You cannot use raw mode if you have configured shared storage in Active Circle. Raw partitions and shared partitions are not compatible, and you may experience problems if both are used at the same time.

To create a partition in Raw mode, you must select a disk with a **Type** of **Generic** (only visible when raw mode is activated) and then click **Add Partition**.

Because the partition will occupy all the space on the disk, you do not need to supply the path or the size of the disk. Just enter a name and click **OK**. The partition appears in the table below the disk with which it is associated.

**Warning**

When operating a partition in Raw mode, the partition table for the disk is over-written. It is therefore important that you prevent these disks from being automatically mounted by the file system. If you do not, the server will not be able to start.

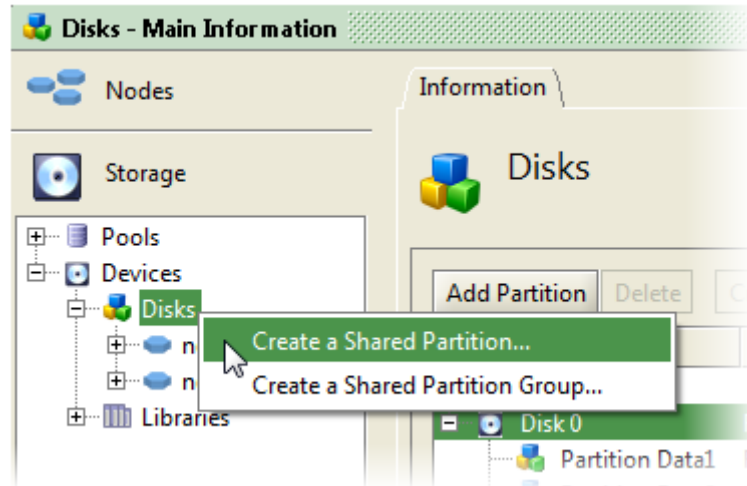
Adding a Shared Partition

To be used in shared storage in Active Circle, a shared partition must be defined by a unique, dedicated LUN (volume label) in the network, using a SAN architecture or by an exported NFS volume. This resource must be configured to be available to all the node servers you want to include in the shared disk pool(s). This may require the restart of the server machines of the nodes. Only one partition should be configured for each LUN. The partition must *not* be mounted locally on any of the servers in the network. Active Circle will automatically elect a

leader node to mount the shared partition on when the shared disk pool is created (for more information, see [Chapter 10, Shared Disk Pools](#)).


The creation of the shared partition in Active Circle can be launched from any node that has access to the label (LUN) of the partition. To create a shared partition, right-click on the **Disks** item under **Devices** in the **Storage** view.

Figure 13.4. Create A Shared Partition



The **Create a Shared Partition** dialog box is displayed.

Figure 13.5. Create A Shared Partition



Fill in the fields to identify the shared partition, following these instructions carefully:

Label

The label should be a unique identifier of the partition in the network. It can be a LUN corresponding to the actual volume label of a partition as defined in a SAN environment. The label can also be specified as an NFS mount point in the format `servername:/path/to/partition`. The resource represented by the label *must* exist and be visible to all the nodes that will be included in the shared pool using the partition. **Important:** Be careful to spell the label exactly like it is defined in the file system where it is

Name	hosted. If you type an incorrect name, the system will accept it, but the partition will not be available. This field will be filled in with the label text by default. You may specify a different name if necessary.
Size (GB) / Use free space	Specify the full size of the partition, in gigabytes. The whole partition must be reserved for Active Circle. The simplest way of doing this, is to check the box Use free space , to let the system determine how much space is available when the volume is mounted and use all of it for the shared partition. The partition should be empty. If it contains any data, it may be unusable for shared disk pools. The minimum allowed size for a partition is 10 GB.

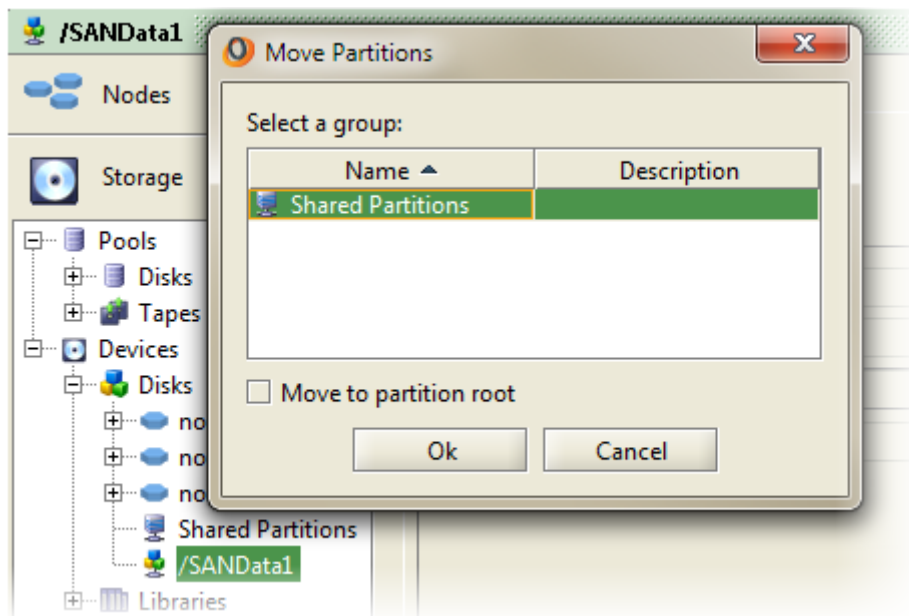
Shared Partition Groups

You may create an element in the tree view under **Devices > Disks** to group your shared partitions. You may then move the shared partition objects to or from one or more of these groups. This will not affect the actual location of the partitions, it is just an aid in keeping the user interface organized.

To create a shared partition group, right-click on the **Devices > Disks** item and select **Create a Shared Partition Group**. Type a name for the group and click **OK**. The group object will appear above your shared partitions in the **Storage** view.

To move a shared partition to a group, right-click on the partition object and select **Move Partition....**

Figure 13.6. Move a Shared Partition to a Group



Select the group you want to move the partition to, and click on **OK**. To move the partition *out* of a group, check the **Move to partition root** option, and the partition will be removed from the group.

Deleting partitions

To delete a *local partition*, select either the **Disks** item or the object representing the node the partition is located on, under **Devices**, and select the partition from the list that is shown in the right pane. If you have selected a partition without pool association (no entry in the **Disk Pool** column), the **Delete** button will be activated. Click in the **Delete** button to delete the partition. You cannot delete a partition that is associated with a disk pool (you must first remove the partition from the pool).

To delete a *shared partition*, select the partition in the tree view under **Devices > Disks**, and right-click on the partition object. Then select **Delete** from the context menu.



Warning

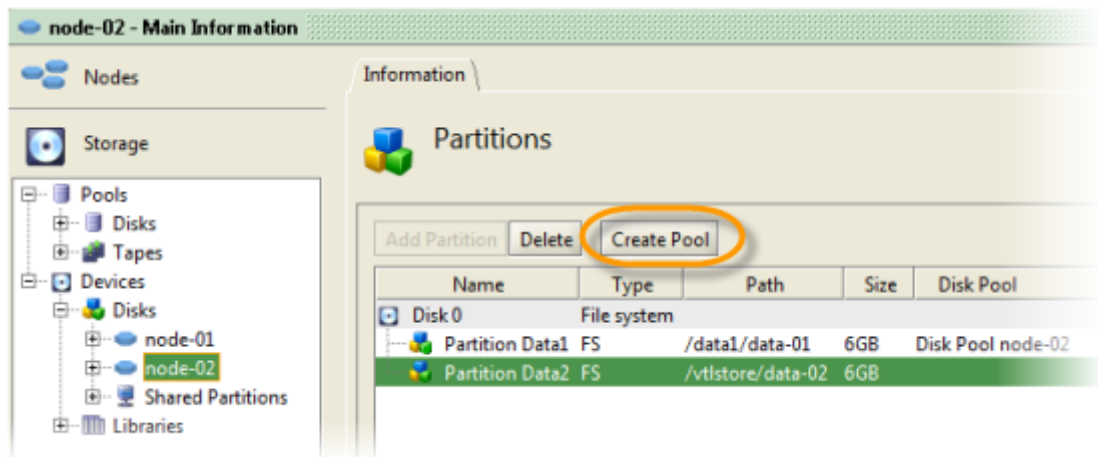
All data on a partition will be lost when you delete it from Active Circle.

Create Pool from Partition

The primary procedure for creating a local disk pool is to use the **Disks** element under the **Pools** item in the **Storage** view. This is described in [“Creating a Local Disk Pool”](#). However, you can also quickly create a pool under **Devices > Disks**, using the list of disks and partitions described in this chapter.

Before you can create a new disk pool, there must be a local partition available that is *not* associated with a pool. Click on either the **Disks** item or the object representing the node the partition is located on, under **Devices**, and select the partition from the list that is shown in the right pane. If you have selected a partition without pool association (no entry in the **Disk Pool** column), the **Create Pool** button will be activated. Click on the button to start creating the pool.

Figure 13.7. Create Pool from Partition



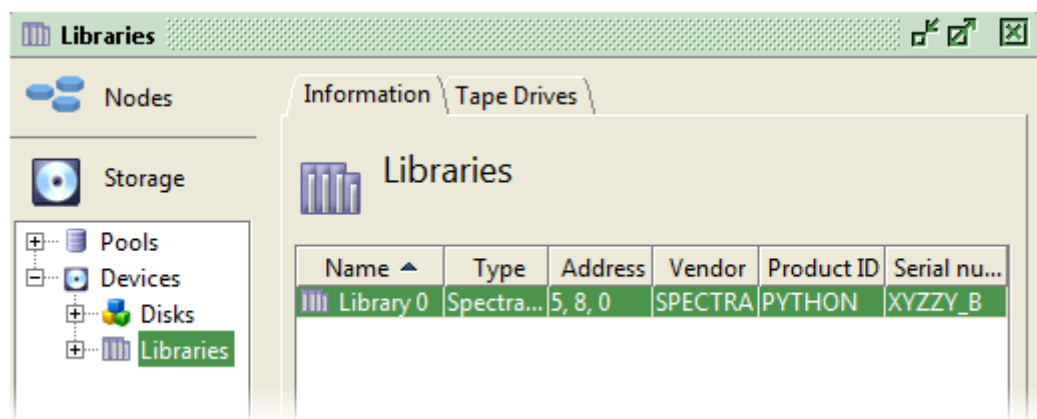
In the dialog which is displayed, enter the name of the disk pool and specify if you want this to be a *cache pool*. Please refer to [Chapter 9, Local Disk Pools](#) for more information on how to configure a local disk pool.

Chapter 14. Library Devices

The Active Circle system allows you to manage most aspects of any tape libraries attached to your nodes. When a node with a supported tape library is joined to your Circle, the software automatically recognizes the library, its tape drive(s), picker, mailboxes, and slots.

You access the library information and operations by opening the **Libraries** item under **Devices** in the **Storage** view. This item (and its subitems) gives you an overview of the libraries of the system as well as allowing you to perform library management tasks on the library **Inventory**, **Tape Drives** and **Nodes** associated with the libraries.

Figure 14.1. Libraries item



Library Detection and Information

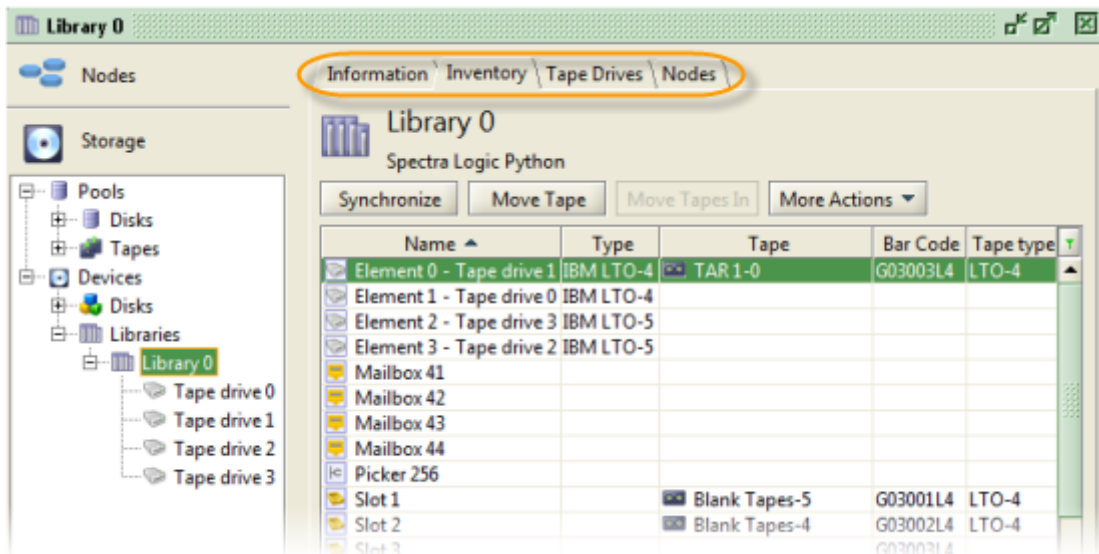
Detection of attached libraries and drives is activated by default.

A properly connected library should appear in the Administration Tool (node restart required if service was started before the library was connected). The library item (consisting of the word "Library" followed by a number, for example "**Library 0**") should appear in the **Storage** view under the **Devices > Libraries** item. In the tree view, any attached tape drives will display under the library item. The library will automatically be shared if it is connected to more than one node through a SAN architecture.

If your tape library is not showing up in the administration interface, check that your tape library is properly connected and configured. For more information on SCSI configurations for tape libraries, see [Appendix B, Configuring and Testing Drives and Libraries](#). You can also verify the detection parameters by checking the node properties in the file `/activecircle/.localvars`. For more information about this file and the SCSI detection parameters it defines, see [Appendix A, Local Variables](#).

When clicking on the library in the tree view, four tabs will be displayed in the pane to the right. Use these tabs to check the information and perform operations on any libraries detected by the node.

Figure 14.2. Library tabs



Information This tab shows overview information for the tape library and its contents. The library name, address, vendor, product ID and serial number are listed. The **Library elements** area shows counters for the number of drives, mailboxes, total slots and free slots in the library. You cannot change any data using this tab.

The **Address** information for a selected library (or drive) displays the device SCSI addresses configured for each node that is connected to it, in the form: *nodename=x,y,z* where "x,y,z" are integers representing the logical unit number (LUN) of the library on *nodename*.

Inventory This tab shows all the contents of the library. Here you can verify that all the elements of the library have been correctly detected by the system. The tape drives, mailboxes, picker, slots and all the tape cartridges should appear in a list in this view, with these information columns:

Name	Bar code
Type	Tape type
Tape	

The table allows you to quickly see the contents of each slot or element, and which tape pool (if any) each tape has been assigned to. Verify that all of the tape cartridges have been detected and that each cartridge's bar code has been properly read. If this is not the case, verify the library, making sure the bar code labels are correctly attached, that the picker is activated, and so on.


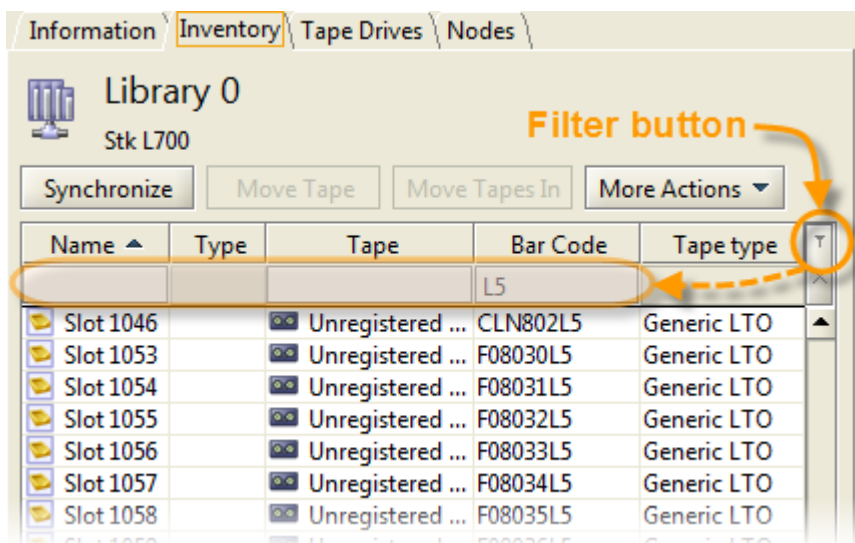
You can apply filters to the table to find items in the library that correspond to element names, tape names or tape bar codes. Click on the filter  icon in the top right corner of the table, and type in characters in one or more of the filter fields to match any of the table content. The list of elements will be updated to match the criteria as you are typing. In the below example, we have typed in "L5" in the bar code filter field to find all bar codes containing those characters.

Figure 14.3. Filters on Library – Inventory tab



This tab also has buttons that allow you to perform operations on the library. For more information, see [Library Operations](#).

Tape Drives

This tab shows a list of all the tape drives in the library. The columns display detailed data for each drive:

Name	Serial number
Type	Number of loads
SCSI Address	Volume read
Library name	Volume written
Tape name	Recovered errors
Vendor	Unrecovered errors
Product ID	

This tab also has buttons that allow you to perform operations on the tape drives. For more information, see [Tape Drive Operations](#).

Nodes

This tab shows which nodes have access to the library, the state and the SCSI address of the library on each node. For more information, see [Access Nodes](#).



Note

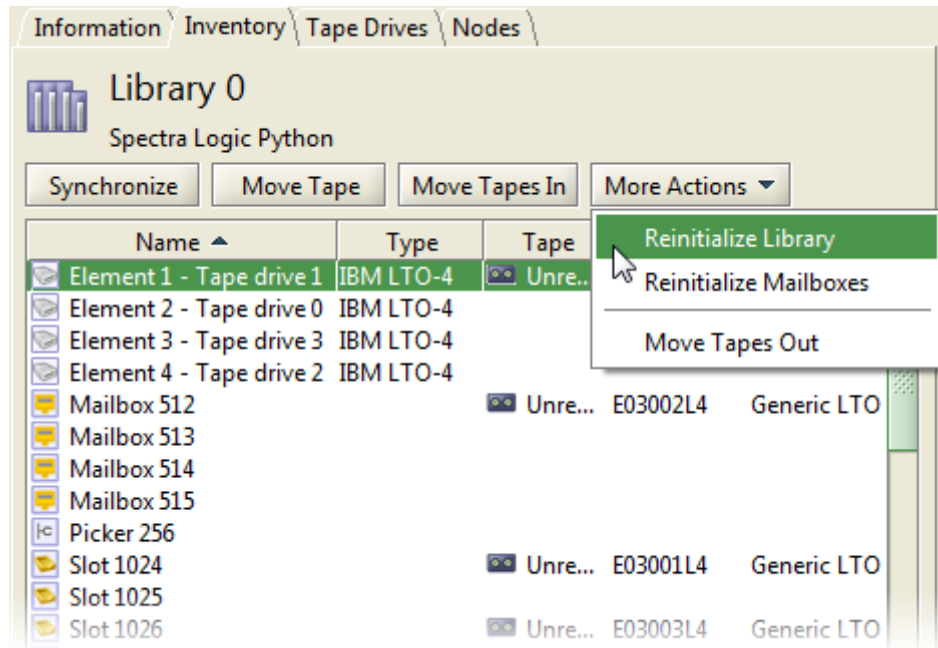
Tapes in tape drives are automatically unloaded and placed in an available library slot after a period of inactivity of 1 hour. This has been implemented to minimize the risk of errors due to hardware or power failures.

Library Operations

To perform operations on the library elements, such as moving tapes, reinitializing the library and mailboxes and synchronizing the library, select the library from the **Storage** tree view **Devices > Libraries** item. Go to the **Inventory** tab, where you can access all the library elements and perform actions on them.

Above the table of library elements is a row of buttons for performing library management tasks.

Figure 14.4. Library – Buttons on Inventory tab



Synchronizing a Library

Sometimes you need to synchronize the display in the Administration Tool with the inventory on the tape library, for example when a new tape has been loaded in the library, or a tape has been removed from the library. To do this, click on the **Synchronize** button. This button acts like a refresh button, querying the library device inventory and updating the display accordingly. Note that this will not perform a re-scan of the slots in the library.

Reinitializing a Library

You can update the inventory of the library or the mailboxes by clicking on the appropriate menu option that is shown when you click on the **More Actions** button.

Reinitialize Library

This option instructs the library device to perform a complete re-scan of all its slots and to update the inventory accordingly. Note that, depending on the number of slots in the library, this operation can take a considerable amount of time and that during the re-initialization none of the data in the library will be available.

Reinitialize Mailboxes

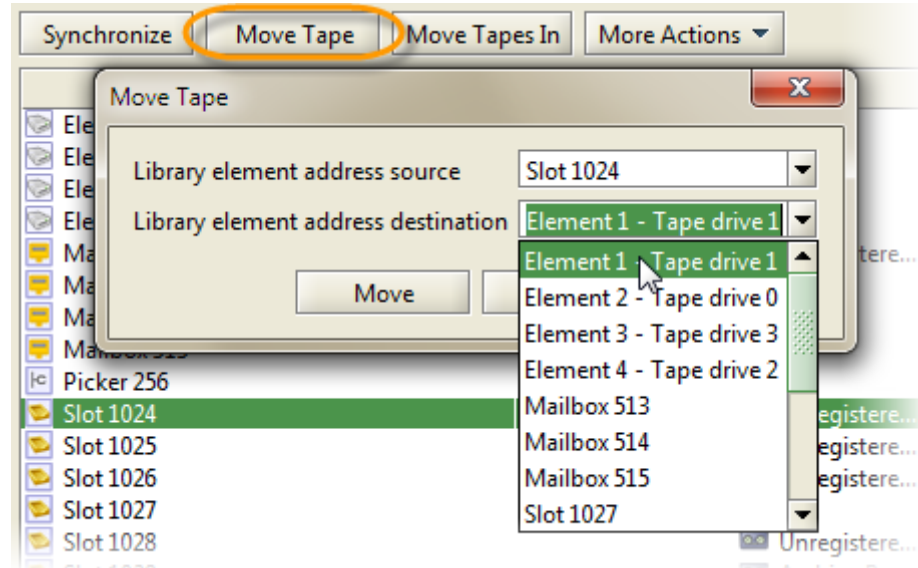
This option sends a command to the library device to scan the Mailbox slots for any changes and update the inventory accordingly.

Moving Tapes

Moving tapes within a library can be done in several ways. The **Move Tape** button can be used to move any tape in the library to and from any slot or element. Changes made here

also perform the corresponding physical move of a cartridge from one slot to another in the library hardware. When you click on **Move Tape**, a dialog box is shown where you can select the source and the destination for the tape.

Figure 14.5. Selecting Source and Destination for Tape



If you selected a slot or an element that contained a tape, the element will be preselected in the **Library element address source** field. If you instead selected an empty slot or element, this element will be preselected in the **Library element address destination**. You can change the selections by clicking on the arrow to the right of each field and selecting an element from the list:

- **Library element address source:** Displays all slots and elements which have a tape loaded in it, including mailboxes
- **Library element address destination:** Displays all empty slots and elements, including mailboxes



Note

If a move is performed by an operator working directly on the library's control panel, Active Circle will not see it until the **Synchronize** button is pressed.

Moving Tapes To and From Mailboxes

In addition to the general move button described above, there are two other options that can be used as "shortcuts" to move tapes in and out of the library's mailboxes. The mailboxes are used by the operator for physically inserting and removing tapes into and from the library. To remove a tape from a library, the tape must first be moved to a mailbox. To load a tape into a library, it must first be placed in a mailbox. You can move tapes between the mailboxes and the other slots in the library, using two special options:

Move Tapes In

This button allows you to move any tapes located in the library mailboxes into empty slots in the library. The system will automatically select the first available slot. All the tapes in the

Move Tapes Out

mailboxes will be moved. The button is greyed out if there are no tapes in the mailboxes.

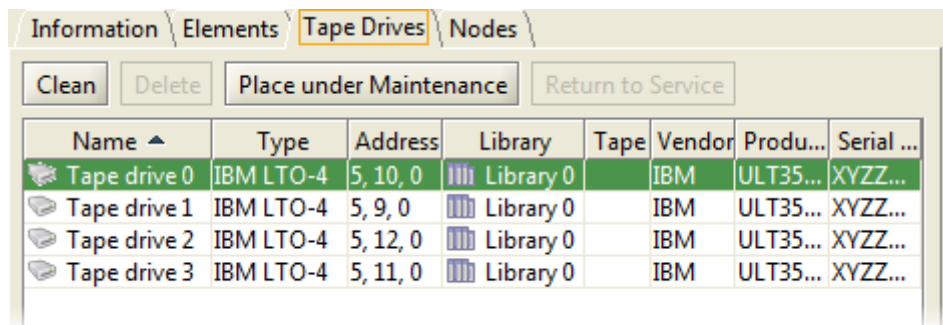
This option is accessible as a menu item when clicking on the **More Actions** button. It allows you to move the selected tapes to mailboxes, for the purpose of removing them from the library. You can select multiple tapes by holding down the **Ctrl** key while clicking. If you select more tapes than the number of mailboxes available, no tapes will be moved. You will receive an error message, and a supervision note will be issued for each tape that could not be moved. The menu item will be greyed out if no elements with tapes are selected. When moving tapes to mailboxes and subsequently removing them from the library, the tapes will remain assigned to their pools, but they will have the status *external*.

Tape Drive Operations

To perform operations on tape drives, such as cleaning, deleting and placing under maintenance, you must first select the library the tapes belong to or the general **Libraries** item under the **Storage** tree view **Devices** item. Go to the **Tape Drives** tab, where you can access all the tape drives and perform actions on them.

Above the table of tape drives there is a row of buttons for performing tape drive tasks.

Figure 14.6. Library – Tape Drives tab



Place Under Maintenance

Use the **Place under Maintenance** button to take a tape drive off-line. A wrench will be added to the drive icon when it is under maintenance.

Figure 14.7. Tape Drive Under Maintenance



By placing a hardware element under maintenance, you are essentially making it unavailable to users. This can be useful when an operator needs to perform an operation on a drive.

When you want to make the drive available again, select it and click on the **Return to Service** button.

Deleting a Tape Drive

You can delete a tape drive from the Active Circle catalog if the drive has already been disconnected at the library level. Use the **Delete** button for this purpose. The button will be greyed out as long as no disconnected tape drive is selected. A tape drive is disconnected when a red cross is shown on its icon.

Cleaning a Tape Drive

Tape drives need to be cleaned on a regular basis, the frequency of the cleaning depending on the technology of the drive and the amount of use. This section describes the general procedures required to initiate a tape drive cleaning process for libraries with cleaning media.



Note

During cleaning, the tape drive will be put in maintenance mode.

Cleaning Notification

When a tape drive detects that it needs cleaning, it emits a warning that is caught by Active Circle. There are two types of notification possible:

Preventive Notification

Active Circle regularly queries the tape drives to check whether they need cleaning, for example when mounting a tape. If the drive responds that it needs cleaning, a supervision note is sent to the administrator. As soon as the drive is no longer reserved for reading or writing, the cleaning begins automatically.

Corrective Notification

In some cases, it is not possible to respond to a preventive notification. The drive might not support sending messages or it might not have been sent. After a certain amount of time, the drive will refuse to either read or write and a supervision note is sent to the administrator. When this occurs, you should immediately suspend all activity on the disk and run a manual cleaning operation on the tape drive.

The tape cleaning procedure generates a supervision note.

Automatic Cleaning

Automatic cleaning is the normal state of affairs and requires no intervention on the part of an operator. When a drive is dirty, Active Circle detects it, loads a cleaning tape in the drive, and the cleaning takes place. When finished, the cleaning tape is automatically removed and placed in an empty slot. If the process is successful, the associated supervision note is automatically closed and the drive is again available for reading and writing.

Manual Cleaning

A manual cleaning operation can be initiated at any time. To launch a cleaning process, go to the **Storage** view in the Administration Tool, click on the **Storage** tab and then the **Tape Drives** sub-tab.

Select a tape drive to clean and then click the **Clean** button.



Note

The system does not ask for confirmation before beginning the cleaning. However, most tape drives refuse to be cleaned if they are not dirty. You must also make sure there is a cleaning tape assigned to the **Cleaning Tapes** pool.

A supervision note is sent to the administrator with information that a manual cleaning operation has been initiated.

Events During Cleaning

Cleaning tape missing or worn out

When the system cannot locate a cleaning tape, or when the only tapes available are worn out, a supervision note is sent to the operator, who must then put a cleaning tape in the appropriate library and then assign it to the **Cleaning Tapes** pool.

Once this is done, the operator can open the note and click on the **Clean** button in the lower left-hand corner of the note.

Restarting during cleaning

If a node is restarted during a cleaning cycle, the drive will complete the cleaning once the node has started back up. The note must be closed manually, however.






Important

If the node restarts before the cleaning tape is mounted in the drive, it will not be mounted after restart. You will have to proceed with a manual cleaning process.

Access Nodes

Use the **Nodes** tab for a library or a tape drive to check which nodes have access to the library or drive.

Figure 14.8. Library / Tape Drive – Nodes tab

Information \ Inventory \ Tape Drives \ Nodes			
Name ▲	Access	State	Address
 node1	Yes	In Service	5,8,0
 node2	No	Off	
 node3	No	Off	

The lists contains all the nodes in the circle, with data columns showing information relating to the library or drive:

Name The name of the node.

Access Indicates whether the node has access to the selected library or tape drive (value is either "Yes" or "No"). For shared libraries, the **Access** column would

display "Yes" for all the nodes which have access to the SAN-configured tape library.

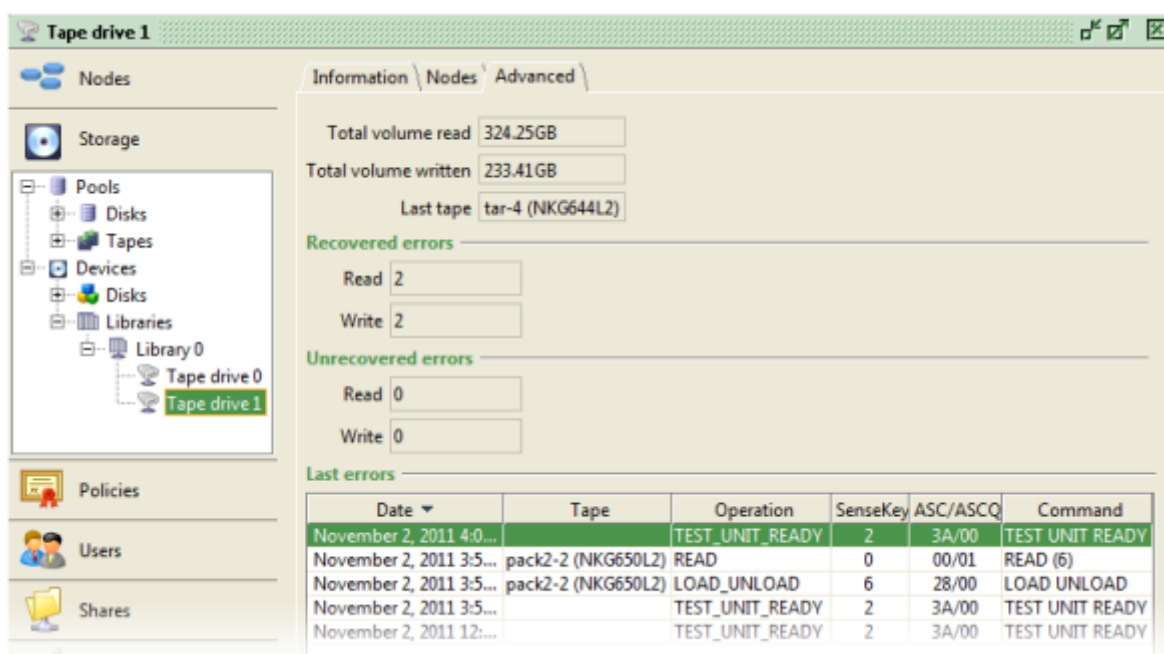
State Indicates if the library or drive is started on the node (value can be "In Service", "Starting", "Off" or "Suspended").

Address Displays the device SCSI addresses configured for each node that is connected to it, in the form: `nodename=x,y,z` where "x,y,z" are integers representing the logical unit number (LUN) of the library on the node.

Advanced Tape Drive Information

The **Advanced** tab for a tape drive shows information about the total volume of data that has been read and written by a drive since it was registered in the Active Circle system, which tape was last loaded and detailed tape error information for the drive. The tape is indicated with both the name and the bar code in the format `tapename (barcode)`.

Figure 14.9. Tape Drive – Advanced tab



The error sections display a summary of the number of recovered (corrected) and unrecovered errors, with separate counts for read and write errors in both categories. The list at the bottom lists details about the last errors that occurred on the selected drive.

By default, the **Last errors** list contains the last 5 errors for the drive. Each entry contains information about the date, the tape used, the operation performed (read/write/positioning), the SCSI sense data in the form of the Sense Key, Additional Sense Code/Additional Sense Code Qualifier (SenseKey, ASC/ASCQ) and the command executed. For more information about SCSI sense data codes, please see the Wikipedia article at: <http://en.wikipedia.org/wiki/ASC/ASCQ>.



Tip

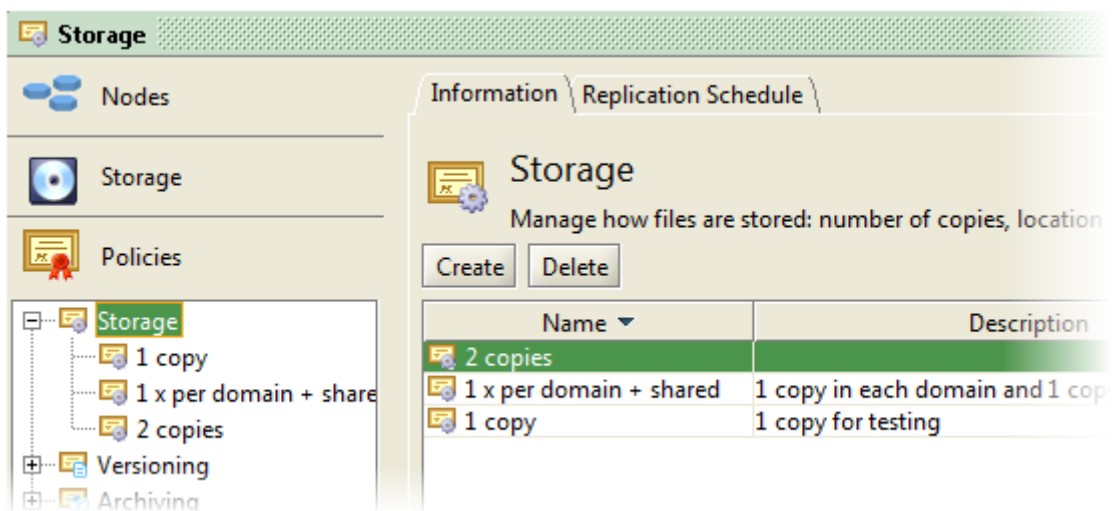
The same type of information can be displayed for a specific tape. For more information, see "Advanced" in [Chapter 11, Tape Pools](#).

Chapter 15. Storage Policy

Storage policies are used for managing how files stored in a share are handled by the system: how many copies of any given file are replicated across the system and where they are replicated, how long the copies should be retained, after how long they should be deleted, if ever, and so on. Note that these rules apply to each version of every file in the system.

The storage policies used for managing how the files stored in the system's shares by users are replicated in the system. By allowing you to configure how many copies of your files you wish to keep, where the additional copies are stored, for how long, and on which media, the Active Circle system enables you to implement a rigorous and economical storage strategy.

Figure 15.1. Storage Policy Overview



To view detailed information about a policy, click on it in the tree view on the left or double-click on the policy name in the list on the **Information** tab.

Replication

When you first create a storage policy, a single basic rule is created whereby each file (or every new version of every file, if version management is activated) affected by this policy will be saved in one copy for an indefinite period of time. That is, the files are not replicated anywhere else in the system. If the disk where those files are stored is damaged, those files would be lost.

In order to provide a minimum of security for your data, you need to modify the constraint in such a way that at least two copies are stored and that they are stored on separate disks, ideally in different physical locations (for example, one file copy on each node).

When files are added to the system, Active Circle first stores them temporarily in the local cache of the node on which the share is mounted. When there is a suitable number of new files, they are then "packed" and moved into the node's pool. This is the default behavior and requires no additional configuration. When the storage policy is executed, it detects that there are new files that require replication in order to satisfy its rules. The files are then copied to the new location (or locations).

Hierarchical Storage

In addition to managing the physical location of the files and their copies, the storage policy allows you to manage your files over time. By configuring complementary storage constraints, you can optimize the utilization of your available space by migrating files to more economical media according to how recently they have been last updated, a strategy known as *Hierarchical Storage Management*, or HSM.

For example, you can configure a storage constraint that keeps files on a fast disk close to where the file owner or users are located during the early, most active period of the file lifecycle, thereby providing optimum performance for file access. Another storage constraint (or rule) could specify that after a period of inactivity, a year for example, the file should be moved to a location that provides high-capacity storage, though at the expense of reduced performance. The constraint might be configured to keep the file in this location for another year, at the end of which it would be migrated by a third storage constraint to a tape library.

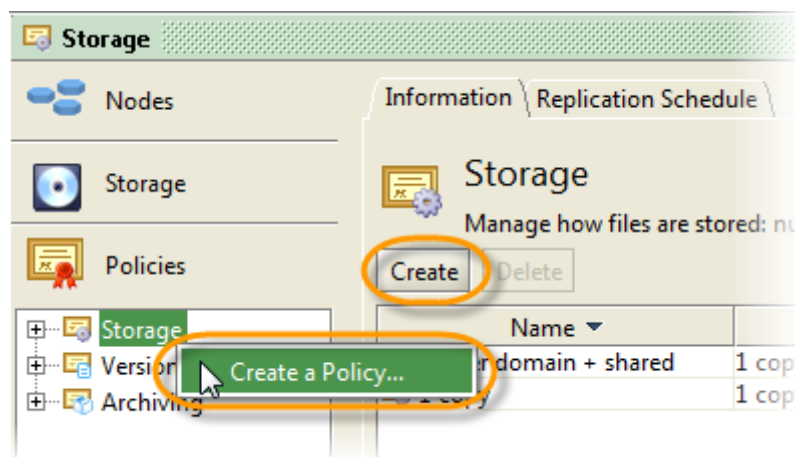
After an appropriate time period, the file could then be archived to tape through the use of another policy: the Archiving Policy (see [Chapter 17, Archiving Policy](#)).

The whole process is transparent to the end user, with processing happening in the background while the file appears to be located in a directory on a network share the whole time.

Create a Storage Policy


To create a storage policy, right-click on **Storage** in the **Policies view** and select **Create a Policy....** Alternatively, you can select the **Storage** item and then click on the **Create** button on the **Information** tab.

Figure 15.2. Creating a Storage Policy



A dialog is displayed where you can enter the policy name and the number of file copies the policy should create on disk or tape.

Figure 15.3. Policy Creation dialog

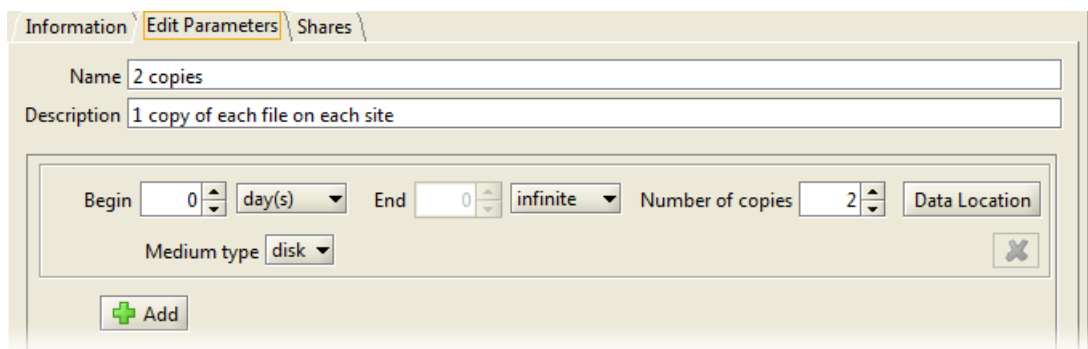


After you click on **OK**, the newly created policy appears as a new entry under **Storage** in the **Policies** view tree structure. The policy will be created with the number of copies you specified and with an unlimited duration. You can change these settings and add more rules for the policy on the [Edit Parameters](#) tab.

Configuring Storage Constraints

Use the **Edit Parameters** tab to configure how and where the storage policy stores your files.

Figure 15.4. Storage Policy – Edit Parameters tab



At the top of the tab are two text fields where you can modify the **Name** and the **Description** for the policy.

On a newly created storage policy, the **Edit Parameters** tab will display a single, basic rule (a set of constraints). This rule tells the system to keep only one copy of each file, for an indefinite length of time. This rule is not sufficient for ensuring replication of your data. You will probably want to configure the policy to provide your organization with the most efficient data storage strategy possible within the limitations of the available hardware. To achieve this, you will need to modify the *storage constraints*.



Note

The policy rules will be specified based on your Circle configuration and the domains, nodes and storage pools defined within it. It is therefore a good idea to install all the nodes and configure the storage hardware that currently makes up your Circle, *before* you start configuring one or more storage policies. Later, you can always add more nodes or storage devices and add them to policy processing.

By configuring the storage constraints, you can specify the duration of the constraint, the number of copies to store, and where to store them. By combining multiple constraints, you can define the storage policies for the entire lifecycle of your data.

The table below describes the parameters for the storage constraints and their usage.

Table 15.1. Storage Constraints Parameters

Parameter	Usage
Begin	From when the constraint should be applied. Zero indicates from the moment the file is first created in the system. Figures are entered as a number of days, weeks, months, years.
End	From when the constraint will no longer be applied. Infinite indicates that the rule never ends.
Number of copies	The number of duplicate copies to store in the system. The value cannot exceed the number of pools available at the data location specified in the rule (for example, you cannot specify 2 copies in a single pool). Copies to be stored in separate physical locations (domain or node) must be specified with data locations in specific pools or with separate rules. See more information below.
Data Location	You can specify a storage location for each policy rule, instead of letting Active Circle decide where to write data.
Medium type	Type of storage device to use: disk or tape.

Defining a Rule

The procedure below provides a step-by-step description of the configuration of the storage constraints.



Tip

Remember that if you make any changes to a parameter, the changes will not take effect until you click the **Apply** button at the bottom of the **Edit Parameters** tab. Before you apply the changes, you can click **Cancel** at any time to delete your changes.

Procedure 15.1. Configuring a Storage Policy

This procedure describes how to configure a storage policy.

1. Duration

First define the duration of the storage constraint by modifying the **Begin** and **End** parameters. In our example, the first storage constraint begins at the moment a file version is created (0 day(s) and ends after 1 month(s).

Information | Edit Parameters | Shares

Name 2 Copies

Description 1 copy of each file on each site

* Begin 0 day(s) End 1 month(s) Num

Medium type disk

+ Add

2. Number of copies

Next you must specify how many copies of each file version should be retained. In our example, we will be storing two copies of each file version for up to one month. However, we want one file on each site, so we specify 1 copy in this part of the rule.

1 month(s) Number of copies 1 Data Location Delete



Important

To achieve an absolute minimum of security, you should save at least two copies of all of your data on two separate disks, ideally in two distinct geographical location.

3. Data Location

Next you can specify where the file version should be stored (this is recommended, but not required) and on which type of medium to store it. To do this, first click the **Data Location** button.

The screenshot shows a configuration window for a storage policy. At the top, there are fields for 'Begin' (0), 'End' (1), a 'day(s)' dropdown, and 'Number of copies' (1). To the right is a 'Data Location' dropdown menu, highlighted with an orange circle and the number 1. Below these is a 'Medium type' dropdown menu set to 'disk', highlighted with an orange circle and the number 2. Underneath is a table titled 'Data Location' with two columns: 'Name' and 'Path'. At the bottom left, there is a '+ Add' button, highlighted with an orange circle and the number 3, and a 'Delete' button to its right.

The **Data Location** table is added below the constraint with columns for **Name**, **Path**, and **Description**. Before you specify the location, you should first select the **Medium type**. This will determine the type of locations that can be added.

4. Medium type

This drop-down menu allows you to choose the type of medium that will be the destination for this particular **Storage Constraint**. The choices are **Disk** (the default) or **Tape**. In our example, we select the medium type **Disk**.


This option can be used, for example, in a context where there is a first rule that states that one copy of the associated files will initially be stored in a particular pool on a fast disk. You might then define a second constraint for the same time period with **Medium type** set to **Tape**. In this scenario, Active Circle will always keep one copy on a fast disk, and one copy on tape in the Active Circle *Pack* format.

5. Add Data Location

Next click the **+ Add** button. This will open a dialog box that presents the storage locations that are available for the medium type you selected. For a **Disk** location, the list of domains, sub-domains, nodes and disk pools in the Circle will be displayed. For a **Tape** location, the list of *Pack* tape pools (pack storage on tape) will be displayed. If no valid storage pools are available for the selected medium type, the dialog box will be empty.

6. Select Data Location

The **Add Locations** dialog box allows you to select the resource for storing the data. If you choose a domain, a sub-domain or a node, the system will choose from among the available pools under the selected object.



The number of resources you select in this list *must* correspond to the number of file copies you specified in the **Number of copies** field for this rule. In our example we want only one copy on the first location, so we choose one resource, in this case the first sub-domain which represents a specific physical site. If you want the files to be stored on a specific node or pool, you can select a resource on that level instead.

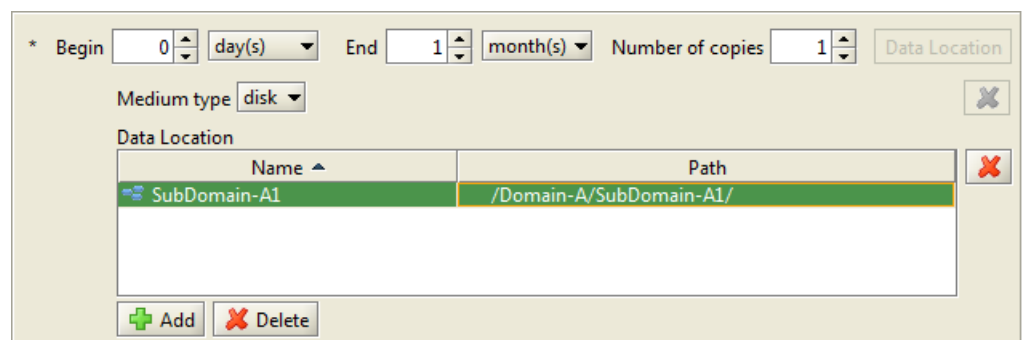


Tip

If you are configuring a rule where the number in the **Number of copies** field is greater than 1, you can select multiple location objects using **Ctrl+mouseclick**.

7. Location Added

Once added, the new storage locations are visible in the table.



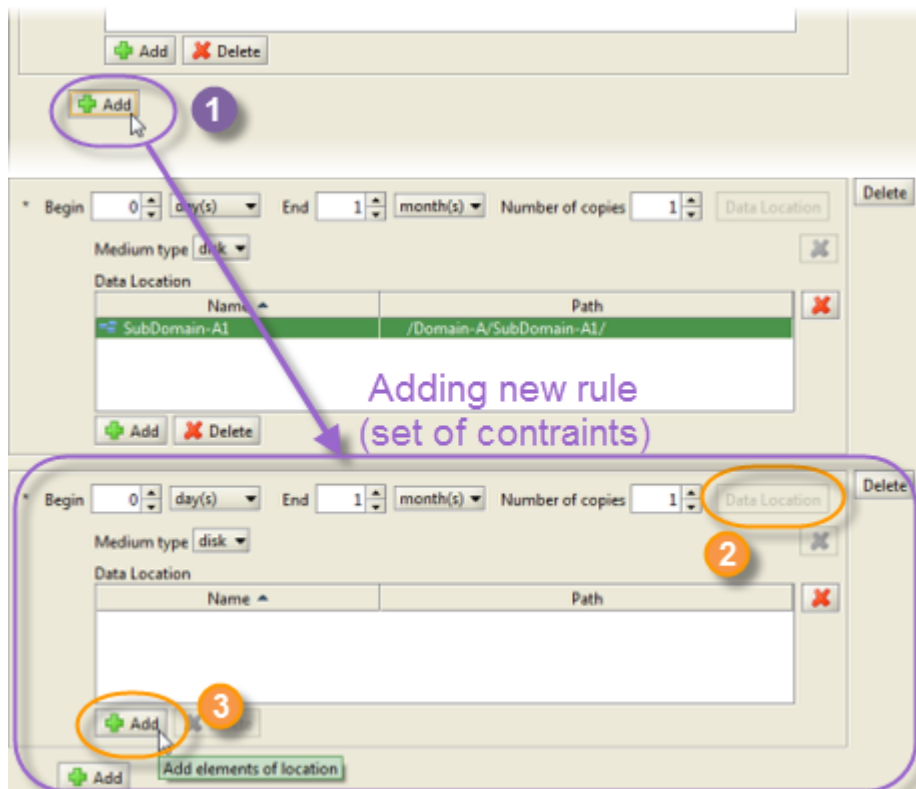


Important

If you specify two or more copies in one single rule, and you do not configure your storage constraint to replicate the files to specific named pools, allowing Active Circle to choose them for you, there is the risk that the system will choose two pools attached to the same node. The system will not write two copies to the same pool, but if more than one pool is available on the same node, Active Circle may write to another pool on that node. If each node has only one pool, the system will automatically write the second copy on another node, as no other pool will be available on the same node. If you want Active Circle to write one or more copies on separate physical locations, regardless of the number of pools on each node, but you want the system to decide on which pools or nodes to write, you *must create a second rule* that specifies the other location. An example of this is shown in the step below.

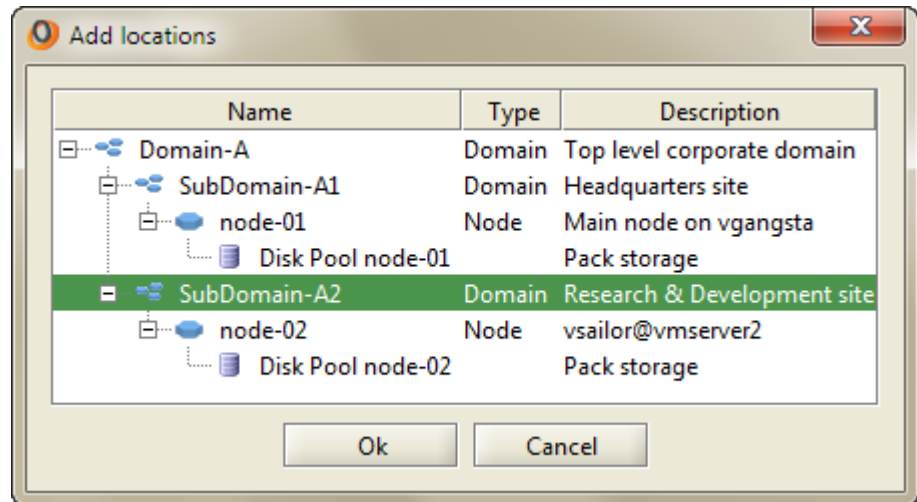
8. Adding Additional Constraints

To add another **Storage Constraint**, click on the **+ Add** button located under the first constraint. This adds another "block" of storage constraint parameters (which make up the rule) that can then be edited to complement the first.



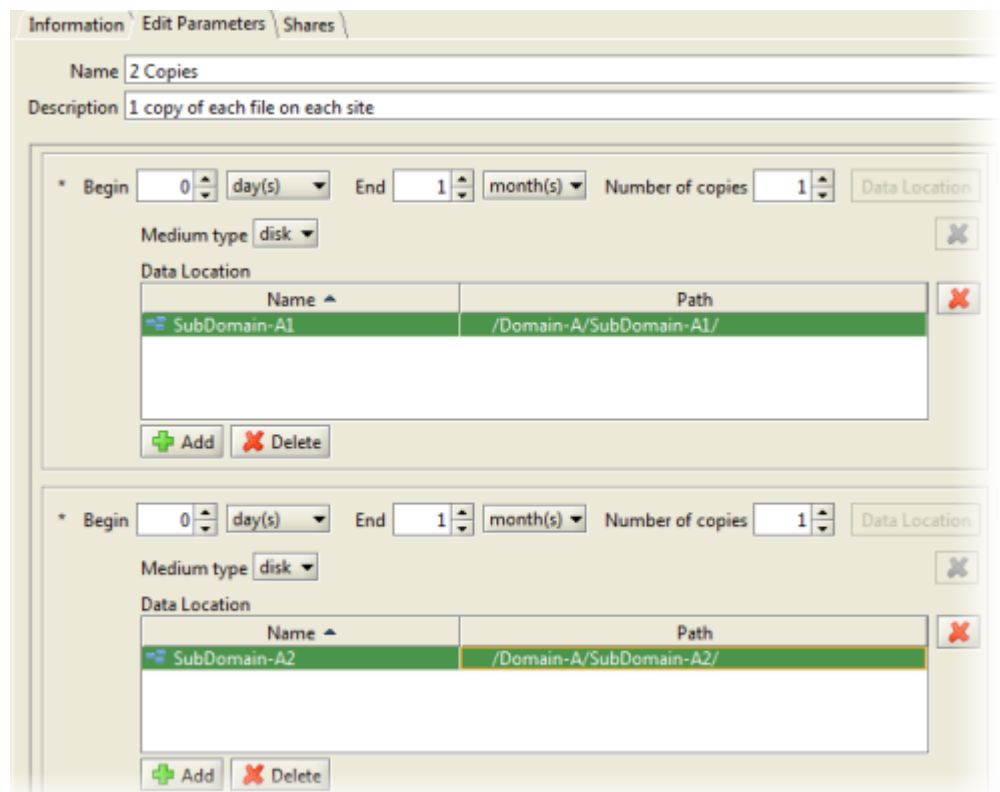
For our example, we want this second rule to be identical to the first rule of the policy, except for the location. So we set **Begin** to **0 days**, **End** to **1 month** and **Number of copies** to **1**. In the **Add locations** dialog box, we select the

other sub-domain of the Circle, to make sure the second copy of the files will be placed there:



We have now created two rules for our storage policy. The first rule ensures that one copy of all files will be stored on disk in the first sub-domain, while the other ensures that another copy will be placed on a disk on a node in the second sub-domain. One copy of each file will then exist in each physical location. The rules will last for one month, after which the replication will end. Another rule or policy (for example a policy to archive the files to tape) should be added to extend the management of the file lifecycle beyond this timeframe. Alternatively, you can set the **End** time for the existing rules to a number of years or infinity.

Figure 15.5. Policy with Two Rules



You can add more rules and different combinations of constraint parameters to achieve the file policy behavior you want. The policy will apply to all files of the shares that the policy is associated with.



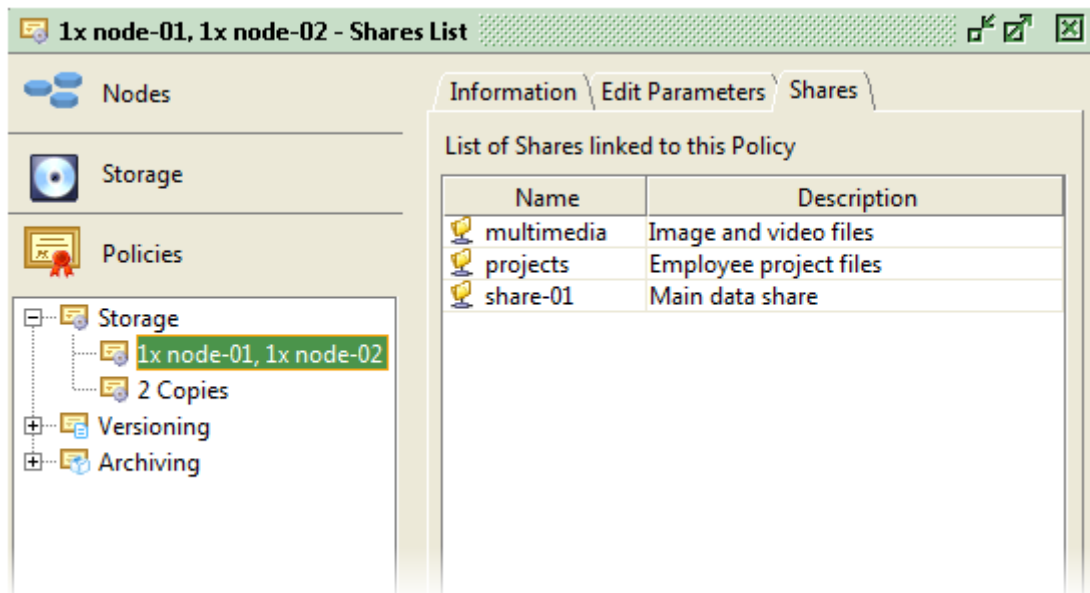
Note

Shared storage pools in a Circle do not belong to any specific domain or node. If you want a shared storage pool to be included in a storage policy rule, you must select the specific shared pool as one of the resources to include, using the **Add locations** dialog box.

Shares Using the Policy

The **Shares** tab displays the shares that are currently using the selected storage policy.

Figure 15.6. Shares Associated with the Policy



For a newly created policy, the list will be empty. Shares are linked to policies during the creation of the share (see [“Create a Share”](#)) or by using the [“Modify Share Parameters”](#) tab for a selected share in the **Shares** view. Only one storage policy (and one archiving policy) can be linked to a share at any given time.

Replication Schedule

Storage policy processing is performed in real time as files are being deposited. Replication between nodes will therefore be done as soon as there are new files to process. For nodes belonging to the same Active Circle domain, this will always be the case.

The storage policy may define that replication should take place between nodes that are located in different Active Circle domains. Depending on the amount of data to be replicated, the process of running the storage policy between domains can consume considerable

bandwidth on a WAN network. This can negatively affect network response times for users. The **Replication Schedule** allows you to limit such *inter-domain* replication by limiting the time periods when the storage policy is allowed to process copies between domains. This functionality applies to all the storage policies defined in the system.

Using this tool, an administrator can easily define a schedule for when inter-domain data replication is permitted, excluding those times of day when the network connecting the domains is most busy. In [Figure 15.7, “Inter-domain Replication Schedule”](#), for example, data replication is not allowed between 8:00 AM and 6:00 PM (18:00), Monday through Friday.



Note

When the system moves from a period during which replication is prohibited into a period during which it is allowed, the storage policies will be automatically run in order to process the files which have not been replicated.

Changing the Replication Schedule

You access the replication schedule by clicking on the **Storage** item in the **Policies** view, and then clicking on the **Replication Schedule** tab (this tab does not exist for individual policies). To make a change to the replication schedule, click on the cell in the table that you want to change: a green cell indicates replication is authorized, a blank (white) cell indicates replication is not authorized. If you make a change, make sure you validate it by clicking the **Apply** button.

Figure 15.7. Inter-domain Replication Schedule

Time	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
0:00 - 1:00	Green	Green	Green	Green	Green	Green	Green
1:00 - 2:00	Green	Green	Green	Green	Green	Green	Green
2:00 - 3:00	Green	Green	Green	Green	Green	Green	Green
3:00 - 4:00	Green	Green	Green	Green	Green	Green	Green
4:00 - 5:00	Green	Green	Green	Green	Green	Green	Green
5:00 - 6:00	Green	Green	Green	Green	Green	Green	Green
6:00 - 7:00	Green	Green	Green	Green	Green	Green	Green
7:00 - 8:00	Green	Green	Green	Green	Green	Green	Green
8:00 - 9:00	White	White	White	White	White	Green	Green
9:00 - 10:00	White	White	White	White	White	Green	Green
10:00 - 11:00	White	White	White	White	White	Green	Green
11:00 - 12:00	White	White	White	White	White	Green	Green
12:00 - 13:00	White	White	White	White	White	Green	Green
13:00 - 14:00	White	White	White	White	White	Green	Green
14:00 - 15:00	White	White	White	White	White	Green	Green
15:00 - 16:00	White	White	White	White	White	Green	Green
16:00 - 17:00	White	White	White	White	White	Green	Green
17:00 - 18:00	White	White	White	White	White	Green	Green
18:00 - 19:00	Green	Green	Green	Green	Green	Green	Green
19:00 - 20:00	Green	Green	Green	Green	Green	Green	Green
20:00 - 21:00	Green	Green	Green	Green	Green	Green	Green
21:00 - 22:00	Green	Green	Green	Green	Green	Green	Green
22:00 - 23:00	Green	Green	Green	Green	Green	Green	Green
23:00 - 24:00	Green	Green	Green	Green	Green	Green	Green

☐ Replication not authorized
 ☒ Replication authorized

Apply Cancel



Important

In order for the replication schedule to be applied to a share, the **Enable data Replication Scheduling** parameter must be explicitly activated in the **Edit Parameters** tab for that share. For more information, please see [“Parameters”](#).

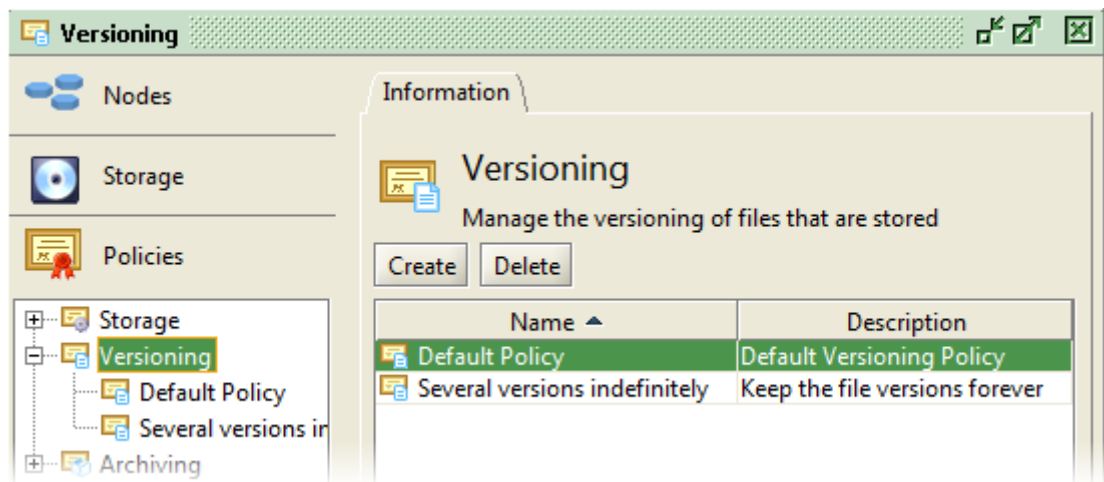
Limiting Inter-domain Policy Processing Load

If defining specific time periods to prohibit inter-domain replication is not practical in a certain environment, an alternative method to limit such processing exists. You can use the *circle property* `cos.slowPackCopyRatio` to define how much of the total policy processing capacity can be used by inter-domain replication. The default value is 100 percent. The suitable value for a specific environment will depend on several factors like expected data replication volume, network load and capacity, etc. For information on how to modify circle properties, see [“Circle Properties”](#) in [Chapter 5, The Circle](#). This value will obviously have no effect during any time periods excluded in the **Replication Schedule** table.

Chapter 16. Versioning Policy

The Version Management (VM) Policy manages the versioning of files that are stored in Active Circle. When activated, the system creates a new copy in the system each time a file is saved by an end user. Although each new file version is linked to the other versions of the same file, these versions are viewed by the system as new files. They are therefore subject to the same rules that are defined in the storage policies, which means that they are replicated across the Circle in exactly the same way. The result is that your files and the history of their development over time are protected by a robust clustered storage architecture.

Figure 16.1. Versioning Policy Overview



To view detailed information about a policy, click on it in the tree view on the left or double-click on the policy name in the list on the **Information** tab.

By default, Active Circle does not permanently remove the files that have been deleted by users, but it removes them from the directory structure that is displayed to the end user. If a user accidentally deletes a file, the file can be recovered using the Active Circle File Explorer.

Active Circle can save multiple copies of all versions of all files, even deleted ones, for an unlimited period of time. Allowing this is probably not a realistic approach as it would require an extremely large quantity of storage space, which works against the purpose of a clustered file storage system. If not properly configured, your available file storage may quickly fill up with multiple versions of your files replicated across multiple sites. Active Circle resolves the problem of preserving the version history of your files without consuming all your storage space, through the use of *sampling*. Instead of storing every version of every file in the system, you may create version management rules that save a sample of all the versions created. Like in the storage policies, you can create a series of rules for a given versioning policy that evolve over time to reflect the probable fact that the older a file version is, the less important it is in the overall lifecycle of the file.

Version Management (VM) Rules

Instead of keeping every version of a file, the system allows you to define sets of complementary rules that will retain a specific number of versions for specific periods of time. For example, a first rule might retain two versions per day for up to one week, after which a second rule might retain two versions per week for a period of six months, and then finally

a rule that states that beyond six months the system will retain one version per month for an indefinite period of time.



Note

If you set the final versioning policy rule to **Infinite**, it means it will be valid for an indefinite period of time. Unless you have an infinite rule, the system will not retain any file versions that are older than the retention period of the last rule (older files will be deleted). This functionality enables the administrator to choose whether to implement a programmed deletion of files of a certain age (last rule defined to one or more years), or to keep older file versions forever (last rule set to infinite).

Default Policy

The **Default Policy** is an empty versioning policy that you can use to facilitate the configuration of the version management process. Because this policy is automatically associated with a share when it is created, you do not need to create a specific policy and then associate it with all the shares you want the policy to apply to. You can just configure the Default Policy.

Predefined rules for the Default Versioning Policy:

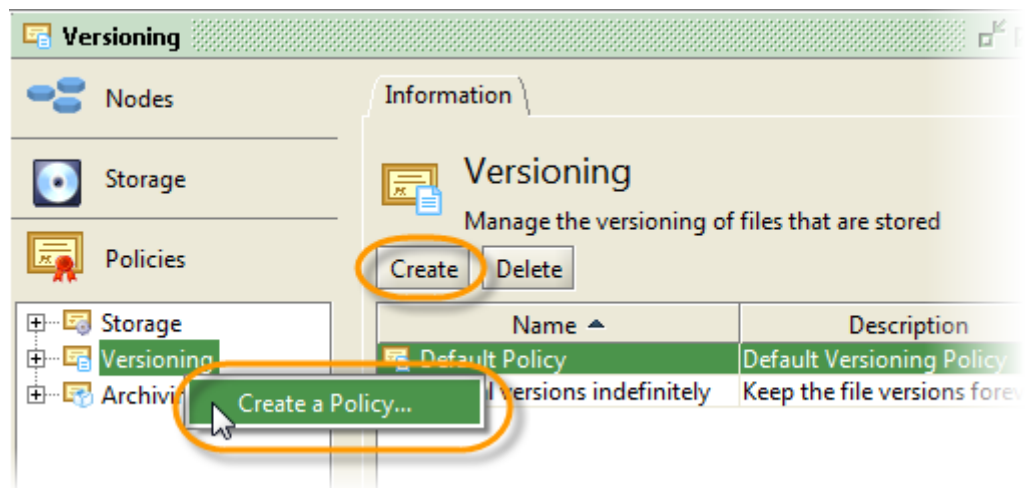
- Keeps one version per day for up to one week
- For file versions older than a week, keeps one per week for up to one month
- For file versions older than a month, keeps one version per month for up to one year
- **NB:** File versions older than one year will *not* be kept (only versions with age less than a year will be retained by the system)
- Deleted files will be kept for 6 months

The Default Policy is automatically associated with shares when they are created. Version management is *enabled* by default, so the system will automatically start storing old versions of your files unless you deactivate version management on the particular share. For information on how to deactivate or activate version management for a share, see [“Version Management”](#).

Create a Versioning Policy


In order to create a Version Management policy, right-click on the versioning policy in the **Policies** view and click **Create a Policy....** Alternatively, you can select the **Storage** item and then click on the Create button on the **Information** tab.

Figure 16.2. Creating a Versioning Policy



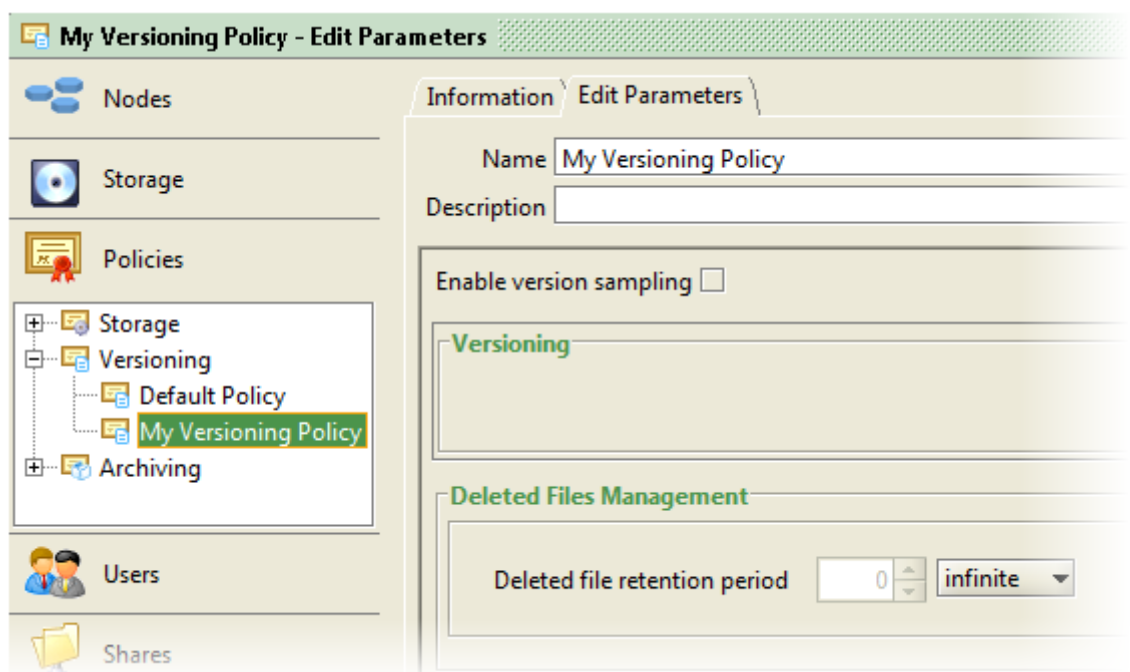
A dialog box is shown to allow you to enter a name for the new policy.

Figure 16.3. Policy Creation dialog box



Type a name and click **OK** to create the policy and it will appear in the tree structure. The **Information** and **Edit Parameters** tabs will show the new, empty policy.

Figure 16.4. Policy Created

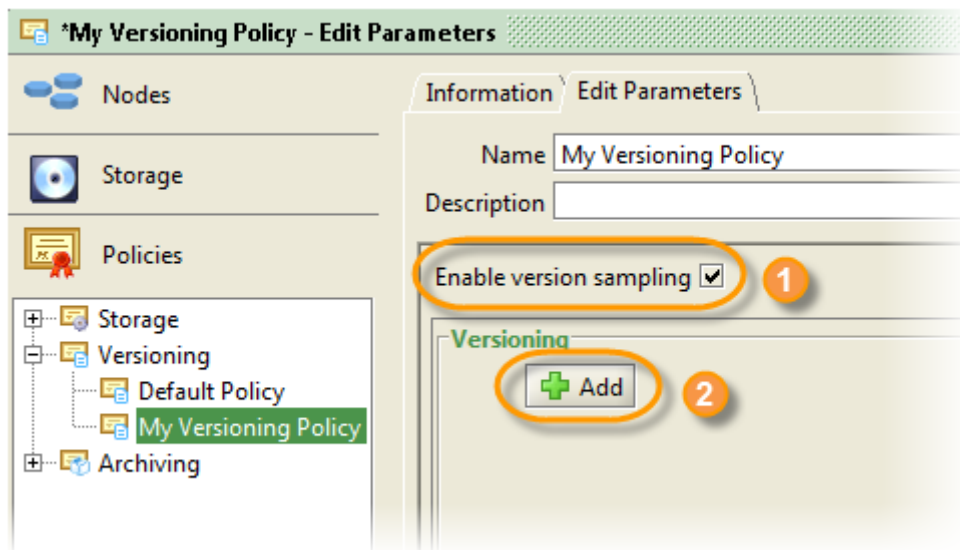


Activation of a versioning policy for a particular share must be done on that share. Please see “[Version Management](#)” in [Chapter 23, Shares](#) for more information.

Adding and Modifying Rules

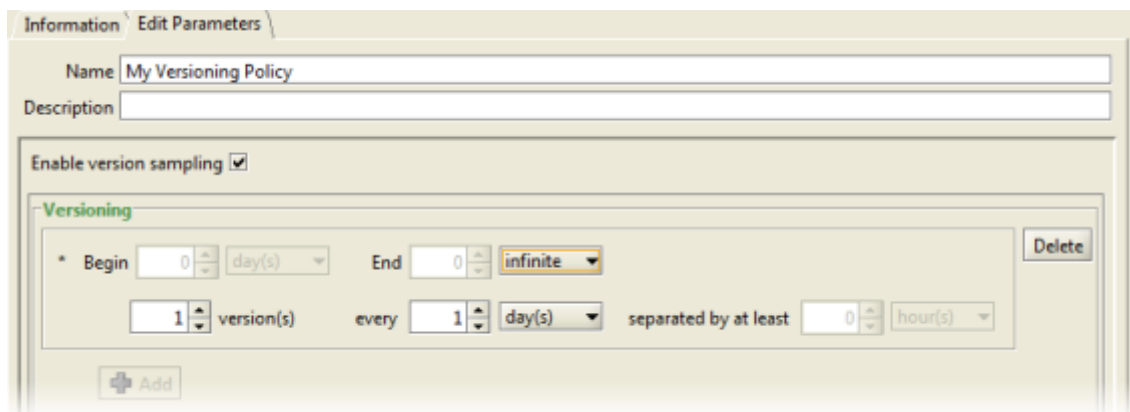
To create the rules for managing file versioning, select a versioning policy and click on the **Edit Parameters** tab. There are two areas in the tab: **Version Management**, where you can create the rules; and **Deleted Files Management**, where you can indicate how long the system should preserve files that have been deleted.

Figure 16.5. Enabling the Policy and Adding a Rule



Activate **Enable Version Sampling** by checking the box. Then click on **+ Add** to create the first rule. The fields for configuring a version management rule will be displayed.

Figure 16.6. Rule Block Added



Adding Rules

You define the policy by specifying the parameters for one or more rules. The parameter fields are described below. Then, to illustrate how the versioning policy works, this will lead you through how to fill in the parameters of a rule by using an example.

Begin	The starting point for the rule. This field is always greyed out because the starting point for any rule is the moment where the last rule leaves off, which for the first rule is the present moment.
End	The end point for the rule. This defines the maximum length of time for which the rule applies to any given file version.
Version(s)	The number of versions of the file to be stored by the service during the period covered by the rule. The total number of file versions stored is determined by combining this field with the next field, Every , which expresses the frequency.
Every	As mentioned above, this field expresses the frequency with which file versions are stored during the period.
Separated by at least	The minimum amount of time allowed between two saved versions. The first file retained will be the most recently created for the period. If more than one file is stipulated for the period, this parameter is enabled to allow you to set the <i>minimum</i> time that should separate any two versions saved. It would not make much sense to take the second-most recently created as this file is likely to be the one that most resembles the first, and is therefore the least representative as a sample. This field is greyed out if you are only saving 1 file for the period.

Example Rules Configuration

This example will configure the following rules:

1. Retain 2 versions per day for a period of 1 week
2. Retain 2 versions per week for a period of 6 months
3. Retain 1 version per month for an indefinite period (Infinite)

Example 16.1. Configuring the First Rule

To configure the first rule, enter the values indicated below:

End	For the first rule, this field is set by default to Infinite and the number field is greyed out. Click on the drop-down list and select week(s) . The number field is now available and it should position itself at 1, which is the value we will use.
Version(s)	Click the Up arrow to select 2 (you will see the Separated by at least field become active).
Every	This field should already be set to 1 day. Keep that value. This means our first rule will retain 2 versions every 1 day for 1 week (a total of maximum 14 versions).
Separated by at least	Adjust this value to 4 hours.

Once you have set these values, click the **Apply** button to save your changes. You can click the **Cancel** button to undo your changes.

Example 16.2. Configuring the Second Rule

Click the **+ Add** button to add another set of rule fields to the list. For this rule, we are going to define the following parameters:

- Beginning after 1 week (this is done automatically)
- Ending after 6 months
- Retain 2 versions
- Every 1 week
- Spaced by a minimum of 2 days

Click **Apply** to save your changes.

Example 16.3. Configuring the Third Rule

Configure the final rule using the following parameters:

- Beginning after 6 months
- Ending never - Infinite
- Retain 1 version
- Every 1 month
- The final option is greyed out because there is only 1 version for the designated period.

Once you have defined your three rules, your screen should look something like this:

Figure 16.7. A Sample Versioning Policy


The screenshot shows the 'Edit Parameters' window for a versioning policy named 'My Versioning Policy'. The 'Enable version sampling' checkbox is checked. The 'Versioning' section contains three rules, each with a 'Delete' button to its right. Rule 1: Begin 0 day(s), End 1 week(s), 2 version(s) every 1 day(s), separated by at least 4 hour(s). Rule 2: Begin 1 week(s), End 6 month(s), 2 version(s) every 1 week(s), separated by at least 2 day(s). Rule 3: Begin 6 month(s), End 0 infinite, 1 version(s) every 1 month(s), separated by at least 0 hour(s). The 'Add' button is at the bottom left.

Understanding the Rules

Now that the rules have been defined, we will look at what they mean.

Imagine that a user creates a file and saves it to a network share that is located on an Active Circle Node. That creates the first version of the file. The user then works on it during the day, saving it five more times over the course of the day. As illustrated below, at the end of the day there are a total of six versions of the same file stored in the system.

Figure 16.8. First Versioning Policy Rule, End of Day 1



The first policy rule states that we want to retain 2 copies for every 24-hour period over the course of 1 week. For our sample file, the service will first retain version 6, the most recently created. It will then look for another version to retain. Version 5 is not retained because it does not meet our requirement of being separated from version 6 by more than 4 hours. Version 4 does meet this requirement and is therefore retained. The other files are purged (permanently deleted).




Important

Note that the files purged are not covered by the **Deleted Files Management** rule (see [“Deleted Files Management”](#)). They are eliminated from the system and no longer recoverable.

Looking at the version history for our file at the end of the first week, we might have something like the following:

Figure 16.9. First Versioning Policy Rule, End of Week 1



As stated in our rule, the system has retained two files for every day, with one exception: there is only one file (version 9) for the second day. This can happen in a situation where no second file meets a condition set in the rule. In our example, we stipulated that the second file must be older than the first by at least 4 hours, but if the user only worked on the file for 1 hour, no file would meet that condition and therefore no additional file version would be retained.

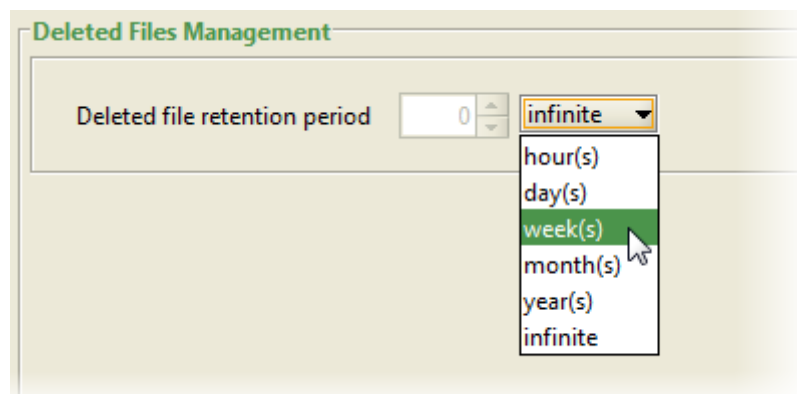
Now let us look at the second policy rule. In it, we stated that we wished to retain two files per week for a period of 6 months. When the policy service is applied to the first week, the versions retained will be version 25, which is the most recent for that week, and version 19, which was created a minimum of 2 days prior to version 25.

The policy service would then continue collecting 2 versions per week (provided that they are created and that they meet the retention criteria) until the 6-month period is completed and the third rule is invoked.

Deleted Files Management

In addition to defining the version management rules, the versioning policy can also be used for defining for how long files that have been deleted by end users should be retained. The default value is **Infinite**; that is, deleted files are retained indefinitely.

Figure 16.10. Deleted File Retention Period



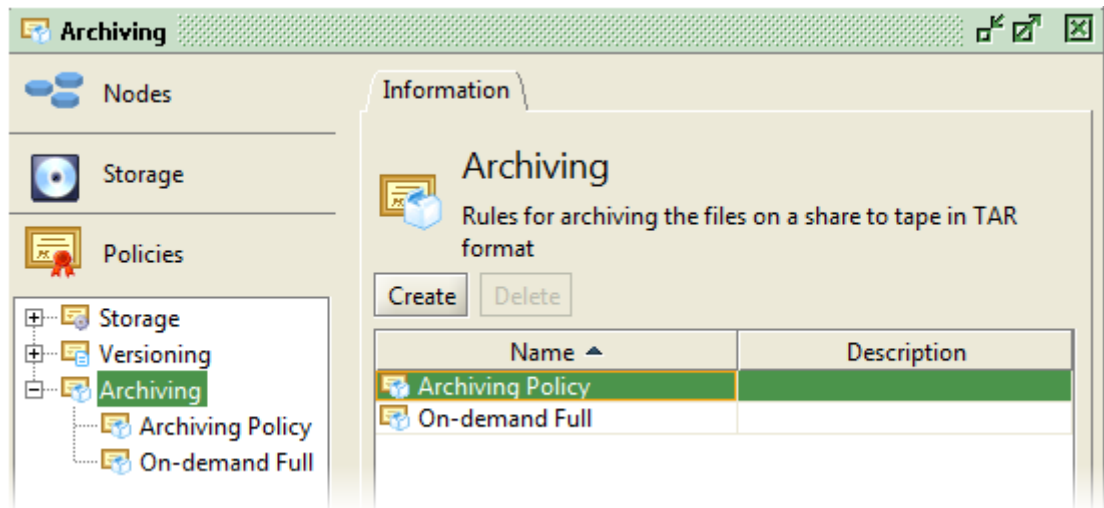
To change the value, select a time frame - Hours, Days, Weeks, Months, or Years - and then enter the number of units.

For example, if you were to choose 6 months, and then a file on one of the shares that are associated with this versioning policy were deleted by an end user, that file and all its associated versions would be recoverable for a period of 6 months from the moment it was deleted.

Chapter 17. Archiving Policy

The Archiving Policy is the service responsible for archiving the files on a share to tape in TAR or LTFS format. Because TAR is an open format standardized in POSIX, its use means that access to files stored in it is completely independent of the Active Circle system. LTFS tapes can also be used outside of Active Circle.

Figure 17.1. Archiving Policy Overview



To view detailed information about a policy, click on it in the tree view on the left or double-click on the policy name in the list on the **Information** tab.

Like the other policies, the archiving policy must be created, configured, and associated with a share. Configuring the archiving policy is almost identical to that of the storage policy, with the exception of two details:

- You can only use tape as the destination medium for an archiving policy.
- Its location must be a previously created **Tape Pool**. You cannot select a domain or a node directly.

The archiving policy is launched in a way similar to the versioning policy. You can configure it to start every day at a certain time or it can be launched manually.

Archiving Properties

Every time the archiving policy is started for a share, a new archive is created which will then be visible in the **Archives** tab for that share. Each archive is made up of one or more TAR files containing the archived files, or an LTFS file structure.

The contents of a single archive are limited to 1 million files by default. Consequently, if the number of files to be archived exceeds this number, more than one archive will be created during the archiving process. The archiving process is incremental, with only those files that have not yet been archived included.

For archiving policies only, you have the option to limit the time allotted to any single archiving process in order to avoid starting a process that might take a very long time (maybe days) to complete, depending on the amount of data to be archived.

The default archiving behaviour and the settings for policy parameters like the ones mentioned above, are defined by *circle properties* of in the `archive` category. These are described in connection with the corresponding input fields under “[Policy Parameters](#)”. If you need an explanation of how to modify circle properties, please see the section “[Circle Properties](#)” in [Chapter 5, The Circle](#).

You can disable the archiving functionality entirely by setting the circle property `archive.disable` to **false**. By default, the archiving functionality is enabled for the whole system.



Tip

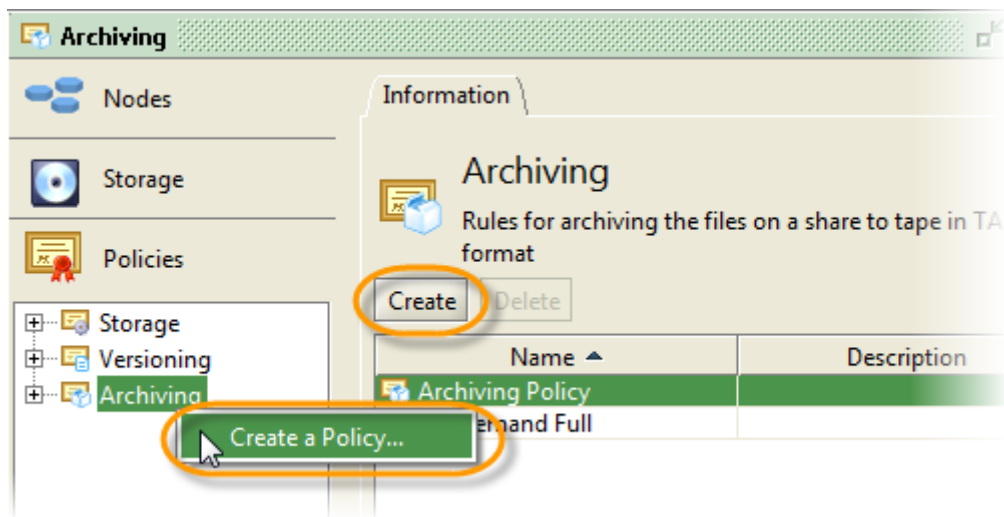
When Active Circle writes an archive to tape, it does so using the `fixed block` mode with a block size of 64 KB. In order to be able to read the TAR file, you must indicate the block size by using the `-b` option to specify the blocking factor (number of 512 byte records). For example:

```
tar -xvbf 128 /dev/nst0
```

Create an Archiving Policy

To create an archiving policy, go to the **Policies** view, right-click on the **Archiving** item in the tree structure, and select **Create a Policy....** Alternatively, you can select the **Archiving** item and then click on the **Create** button on the **Information** tab.

Figure 17.2. Create an Archiving Policy



The Archiving Policy creation dialog box is displayed.

Figure 17.3. Create an Archiving Policy – Name

Archiving Policy creation

Name

Parameters

☒ Data location

Name	Path

☐ Export pool

Type a name for the policy and select the destination for the archive by clicking the **+ Add** button. The **Add locations** dialog box will be shown displaying the current **Tape Pools** available in your Circle.

Figure 17.4. Create an Archiving Policy – Add Location

Add locations

Name	Description
TAR 1	Tape pool on node1
TAR 2	Tape pool in domain B

Select the destination pool or create a new pool by clicking on the **Create a Tape Pool...** button and click **OK** to close the **Add locations** dialog box. Click **OK** again to create the new archiving policy. The new policy will then appear in the policy tree structure and you can review the configuration using the **Information** tab.

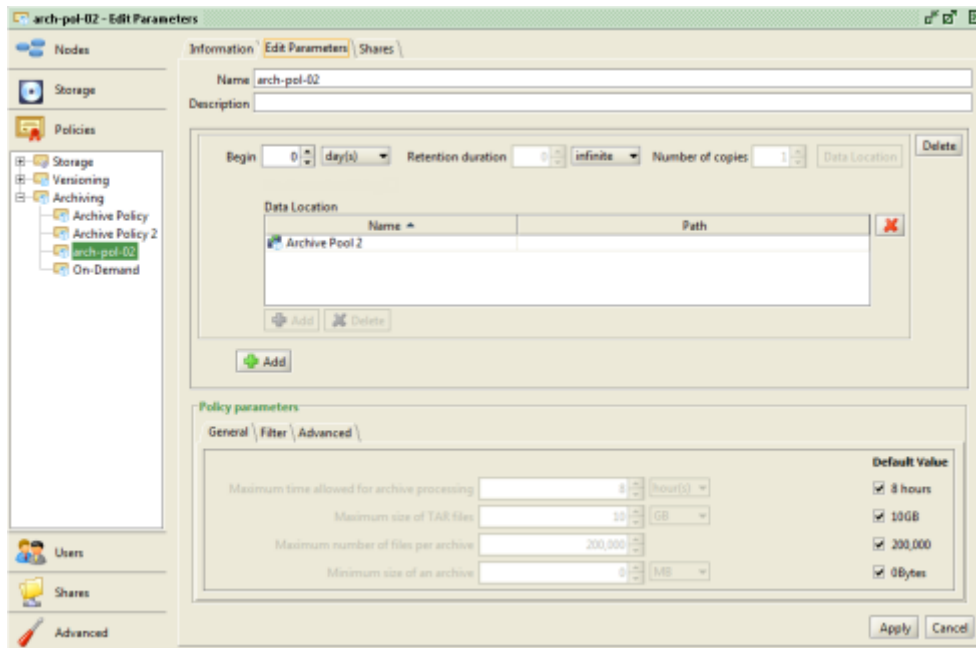


Important

When you are creating a new archiving policy that you want to enable for automatic execution on a share (the normal rule, as described in the above procedure), you must *not* select **Export pool**. **Export pool** can only be selected if an export pool already exists. Export pools are special pools which only allow manual archiving (for more information, see [“Export Pool”](#) in [Chapter 11, Tape Pools](#)).

To configure the newly created policy, select it in the **Policies** view and click on the **Edit Parameters** tab.

Figure 17.5. Archiving Policy Edit Parameters tab



At the top of the tab are two fields: one for the **Name** and another for a **Description**. Below that there are two areas: a set of fields for archiving constraints (explained below) and the **Policy parameters**.

Configuring Archiving Constraints

This area is used for configuring the archiving rules, or constraints. The configuration of the *archiving constraints* is very similar to the process used for configuring storage policy constraints. The primary difference is that archiving rules only apply to tape storage. Consequently, the destination location must be a previously configured *tape pool*. For information on the creation and configuration of tape pools, please refer to [“Creating and Activating a Tape Pool”](#).

You can configure up to 4 rules per policy. The fields for the archiving constraints are as follows:

Begin

The starting point for the archive. When the archiving policy is run (whether manually or as part of a scheduled process), Active Circle will scan the Shares to which the particular policy applies and look for file versions that are at least as old as the parameter given here and which have not been previously archived. All of the file versions that meet these criteria and which do not any criteria for exclusion from archiving will be included in the archive.

Retention duration

The length of time for which the system will retain the archive. In most cases, the default value of *Infinite* would be the proper setting.

Number of copies

Reserved for future use

Data Location

The data location for the first rule is specified during the policy creation. You cannot add another location to the same rule. Use this button and table to specify the location for any additional rules you want to add.

Policy Parameters

You can configure several parameters relating to the archiving policy process. You find these settings in the **Policy Parameters** section at the bottom of the **Edit Parameters** tab for the selected archiving policy. It is divided into three subsections, indicated by separate tabs:

- [General](#)
- [Filter](#)
- [Advanced](#)

All the settings on these three tabs are initially greyed out. The default values for the policy parameters, which are set using the *circle properties*, apply to all the archiving policies. If you want to change any of these parameter settings for a specific archiving policy, you must first uncheck the box(es) to the right of the fields under the heading **Default Value** on each tab. When you deselect the default value box, the parameter field becomes activated, and you can specify another value for the current archiving policy. An asterisk (*) will be displayed to the left of the parameter field until you click on [Apply](#), to indicate that there are unsaved changes pending.



Note

If you want to change the default values which all archiving policies should inherit, you must do so using the *circle properties* indicated for each parameter below. For more information about modifying defaults values, see [“Circle Properties”](#) in [Chapter 5, The Circle](#).

General

Figure 17.6. Policy Parameters – General

			Default Value
* Maximum time allowed for archive processing	8	hour(s)	<input checked="" type="checkbox"/> 8 hours
* Maximum size of archive segments	20	GB	<input checked="" type="checkbox"/> 20GB
* Maximum number of files per archive	1 000 000		<input checked="" type="checkbox"/> 1 000 000
* Minimum size of an archive	0	MB	<input checked="" type="checkbox"/> 0Bytes

Maximum time allowed for archive processing

The purpose of this feature is to prevent the archiving policy from taking over the system during a long

archiving process. For example, you can configure the system to run over night by setting the limit to 8 hours and setting the start time for the service to 10 PM.

Default value: 8 hours

Circle property: `archive.maxDuration`

Maximum size of archive segments

Limit the archived files to a particular size to avoid files of unmanageable size.

Default value: 10 GB

Circle property: `archive.maxTarFileSize`

Maximum number of files per archive

Limit the number of files in a TAR archive. A new TAR file will be created when this number is reached during archiving.

Default value: 1,000,000

Circle property: `archive.maxFilesCount`

Minimum size of an archive

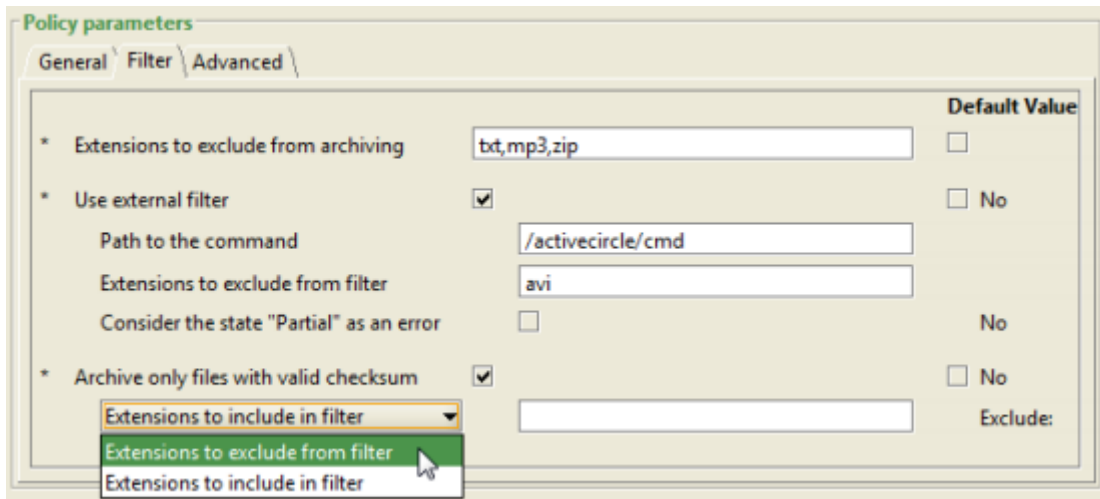
Define a minimum size for TAR files in order to avoid a proliferation of small archives.

Default value: 0 bytes (deactivated)

Circle property: `archive.minSize`

Filter

Figure 17.7. Policy Parameters – Filter



Extensions to exclude from archiving

A comma-separated list of file extensions not to be archived. File extensions should be specified without periods and spaces (for example: `iso,mkv`)

Circle property: `archive.excludedFiles`

Use external filter

To activate this feature, you must first uncheck the **No** check box for accepting the default value to the right of this parameter. Once activated, you can then access the other dependent parameters.

This feature allows administrators to define the path to an external executable that will analyse the data relating to the files to be archived and allows for filtering the results according to specific criteria. The development and the definition of the external executable are the responsibility of the administrator of the Active Circle system.

Default value: No (deactivated). The default value will be changed to "Yes" (activated) if a value is set for the circle property

`archive.externalFilterExecutable.`

Path to the command

The path to the executable to use for filtering.

Circle property:

`archive.externalFilterExecutable`

Extensions to exclude from filter

A comma-separated list of file extensions not to be processed by the filter. File extensions should be specified without periods and spaces (for example: `txt,mp3,zip`).

Circle property:

`archive.externalFilterExcludedFiles`

Consider the state "Partial" as an error

If at the end of an archiving process, one of the archives shows the state "Partial" (that is, that Active Circle was unable to complete the archive), then this should be considered as an error and a supervision note corresponding to this event should be sent.

Default value: false

Circle property:

`archive.externalFilterPartialIsNok`

Archive only files with valid checksum

You can choose to only include files for which the MD5 checksums have been validated, to ensure that no corrupt data will be archived. Note that there is no automatic validation of checksums in the system. You can also enter extensions that you want to *exclude* or *include* in the filter processing.

Default value: false

Circle properties:

`archive.validatedSelection`

`archive.validatedSelectionExcludedFiles`

`archive.validatedSelectionIncludedFiles`



Note

Offline files are excluded from selection by archiving operations.

Advanced

Figure 17.8. Policy Parameters – Advanced

Policy parameters

General \ Filter \ Advanced \

Optimization of the selection ☒ ☐ Yes

Directory level for the organization of archive segments ☐ 0

Keep empty archive ☒ ☐ Yes

Verification after writing ☐

Export Export pool

Optimization of the selection	<p>Enables browsing of the virtual file system in an optimized way for incremental archiving.</p> <p><i>Default value:</i> true ("Yes")</p> <p><i>Circle property:</i> <code>archive.fastTreeWalk</code></p>
Directory level for the organization of archive segments	<p>Close TAR file if folder changes at this depth. Use this parameter to create a TAR file for each folder at a certain level of the folder structure of a share. For example, if the second level in the tree structure on a share contains a folder for each user (/level1/username), specifying "2" will result in one TAR archive file being produced for each user folder (including sub-folders). This archiving mode is also referred to as "structured TAR". A value of zero disables this functionality.</p> <p><i>Default value:</i> 0</p> <p><i>Circle property:</i> <code>archive.closeTAROnFolderChangeDepth</code></p>
Keep empty archives	<p>Keep empty Archives (for which there was not enough data or empty and successful).</p> <p><i>Default value:</i> false ("No")</p> <p><i>Circle property:</i> <code>archive.keepEmptyArchive</code></p>
Verification after writing	<p>This parameter activates a verification of the archived TAR files after they have been written to tape. This verification is not available for LTFS tapes. By default, a checksum is calculated by the system during writing of each TAR file. The Verification after writing option will read all the TAR files for the archive on tape, recalculate the signatures and compare them to the original checksum. This provides a level of protection against any write errors that are not detected by the tape drive. Note that activation of this feature will slow down archiving significantly, since</p>

it implies repositioning of the tape to the beginning of each file after writing it, in order to read the data again. This functionality will not be available if the circle property `archive.disableChecksum` has been set to "true" (default value is "false").

Default value: No

No circle property exists for this parameter.

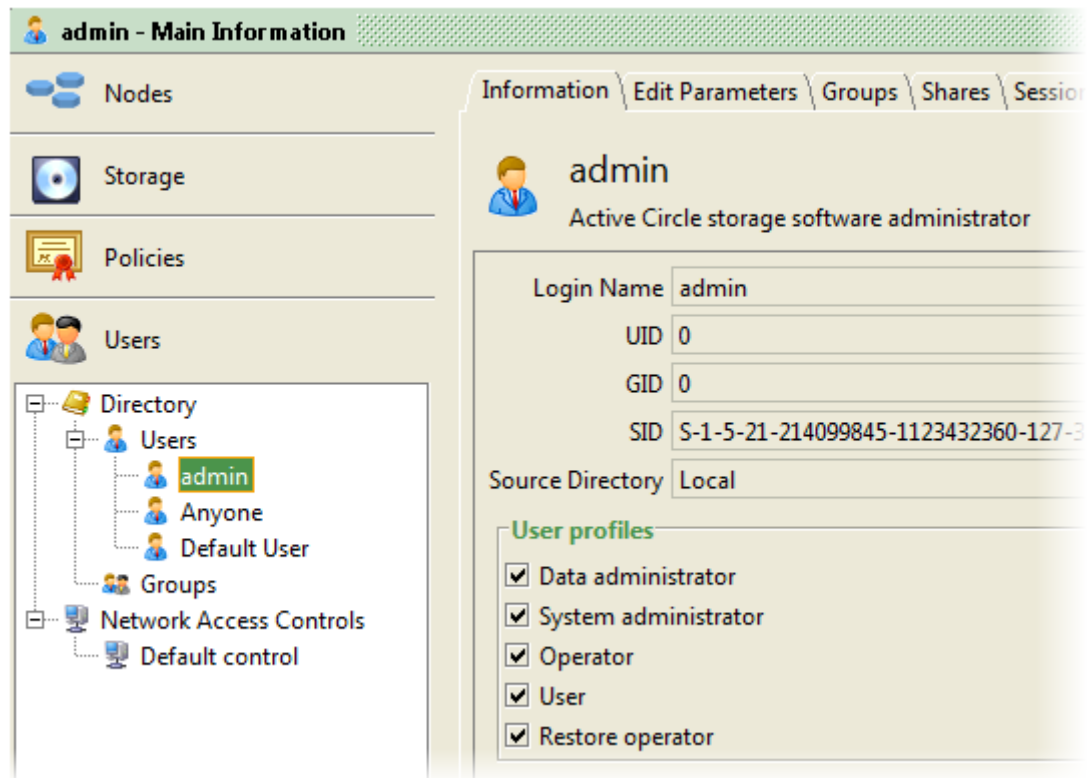
Export pool

Activate or de-activate the [Export pool](#). For this to be available, the *circle property* `archive.enableExport` must be set to "true" and an Export tape pool must exist.

Chapter 18. User Profiles and Access

This chapter provides an overview of the concepts used in managing user and group access in Active Circle. The following chapters will describe the contents of the **Users** view in more detail.

Figure 18.1. Users View and User Profiles



Profiles

Users and groups in Active Circle are assigned one or more **Profiles** that determine what type of actions they are permitted to perform in the system.

Data administrator	Responsible for data storage policies (for example defining policies and creating HA clusters), as well as access to the data (such as assigning user profiles and assigning to groups).
System administrator	Responsible for more traditional IT tasks: installation and management of the system, creation of domains and nodes, managing users, etc.
Operator	Can access all information in the Administration Tool in read-only mode and can launch archiving operations on shares and perform operations on tape libraries, such as closing a tape, moving a tape out, modifying the availability of a tape pool and assigning tapes to pools.
User	To be able to access a share, an account must have the profile named User associated with it. This also applies to accounts with administrative privileges, including the

Restore operator

Active Circle `admin` account. The **User** profile does not have permission to run the Active Circle Administration Tool, but it has access rights for certain operations in the File Explorer application.

Can save local copies of files and directories from Active Circle shares, using the File Explorer application (the **Copy Selection** feature). Other permissions are the same as for **User**. The **Data administrator** profile also includes this permission. *Note that the Restore operator will be able to view ALL files on the shares, regardless of any file permission restrictions that might apply to the user profile. After performing a local copy of the files, the Restore operator will have full rights on the files at the local level.*

Administrator Profiles

The two administrator profiles that are built into the Active Circle system operate on two different levels, and therefore they have different access rights. The System Administrator operates on the Circle configuration level, whereas the Data Administrator operates mainly on the share access level, managing access to the data. Where the system administrator has full access rights, the data administrator has read access, and vice versa. This has been implemented to allow separation of these responsibilities and permissions, so no single profile has all access rights. In many organizations, the two administrator roles would be combined into one. This is the case for the default `admin` user created by the Active Circle system, which has both the administrator profiles. The lists below summarize the rights of the administrator profiles.

SYSTEM ADMINISTRATOR rights:

- Domain and Node management
- User, Group and Network Access Control management
- Directory management

DATA ADMINISTRATOR rights:

- User, Group and Network Access Control *profiles*
- Shares and share access
- Policies
- HA Clusters

For example, the system administrator has the right to create a user, but only the data administrator has the right to change the user profiles for that user. Conversely, the data administrator does not have permission to create users. The separation of rights does not allow the data administrator to create network access controls or import directories, as those tasks are reserved for the system administrator. Once created, the data administrator manages the profiles of these elements.

Directories and Protocols

In order to facilitate the process of user management, the Active Circle Directory can be synchronized with one or more external directories. The types of directories you can import users from, are described in [“External Directories”](#). Once imported, users connect to Active Circle shares using the protocol appropriate to their platform.

The table below lists the supported directory services and the file access protocols that may be used with each.

Table 18.1. Directory Services and Protocols

Directory Services		File Access Protocols		
		CIFS	FTP	NFS
LDAP	Active Directory	✓	✓	--
	OpenLDAP	✓ *	✓	✓
	Sun ONE	--	✓	✓
	Samba	✓	✓	✓
NIS		--	✓	✓
NT4 / Samba (Domain or Server)		✓	✓	--

* Requires the plug-in NTLM ("NT Lan Manager", now known as "Integrated Windows Authentication" or IWA). May also require the recompilation of the OpenLDAP server.



Important

Because of the way in which the Windows challenge method of authenticating a user works, only one directory using this mechanism (Active Directory and OpenLDAP using a Samba schema) may be imported into Active Circle.

Authentication

Active Circle supports different authentication methods. The way in which a user is authenticated depends in the first instance on whether or not the user account is a local one; that is, the user was created directly in the Active Circle. If the user is local, then Active Circle takes responsibility for verifying the user name and password.

If the user is not local, their credentials will be verified by the directory service from which they were imported.

The table below shows the available access protocols and the verifications that can be made when a user connects.

Table 18.2. Protocols and Verifications

Protocol	User Authentication		Network Access Control	
	Login/ Password	UID	IP Address	Host Name *
CIFS (port 139)	✓	--	✓	✓
CIFS (port 445)	✓	--	✓	✓
FTP	✓	--	✓	✓
NFS	--	✓	✓	✓

* Requires correct DNS configuration on all Active Circle nodes and clients. For all protocols, both the short name and the fully qualified domain name are valid for access control by Active Circle.

Access Control Overview

The access to data in Active Circle is managed by several different mechanisms. This section offers an overview of the access control hierarchy and how the access rights, permissions and ACLs are used in combination to grant or limit access to resources.

Data access is managed at three levels that are successively verified. Once authorization is given at one level, the system proceeds to the next level of verification. Access needs to be granted at all levels before the requested operation (for example reading a file) can be carried out. Conversely, if one of the levels denies access, the operation fails. The authorization process follows this order:

1. Verification of [Access Rights on Active Circle Shares](#)
2. Verification of [Unix Permissions](#)
3. Verification of [Access Control Lists \(ACL\)](#)

Access Rights on Shares

The first access control verification is done at the Active Circle share level, where access is verified both according to *user and group access rights* and by *Network Access Controls* (using a hostname, IP address or IP address range).

The **Access Rights** in Active Circle determine what actions a user is allowed to perform on the folders and files on the share level. How to set and modify these access rights is described under "[Share Access](#)" in [Chapter 21, Users and Groups](#).

When a user logs in to the system, the privileges and access rights granted on the user and group level are compared with those granted to the machine as defined in the **Network Access Controls** (for more information on how to define and apply these, please see [Chapter 22, Network Access Controls](#)). Only those rights that exist in both will be retained. For example, if a user only has the "User" profile, while the Default Network Access Control granting the profiles for "Data Administrator", "System Administrator", and "Operator" has been applied to the share, then the user will still only have user-level access to the share.

This first level of control cannot be overridden by modifying ACLs or Unix permissions on files or folders. For example, this makes it possible for an administrator to limit the access of selected users to "read-only", even if the ACLs would have allowed more access to the same users. These rights are computed when the user session is opened.

Unix Permissions

The second level of control is the verification of Unix-like permissions on the files or folders. The Unix permissions will be verified regardless of which protocol is used for accessing the data. These permissions have three modes: Read (r), Write (w) and Execute (x).

By default, the Unix permissions are defined as follows:

- When a file is deposited using the CIFS protocol (Windows), the permissions will be set to full access for everyone: `rwxxrwxrwx`. The default setting can be modified on the node, domain or circle level, using these *Circle Properties*:

```
nas.directoryDefaultMode
nas.fileDefaultMode
```

For more information on how to modify values for circle properties, see the topic "[Circle Properties](#)".

- When a file is deposited using the NFS protocol, the Unix-like permissions will be set to the current permissions on the NFS client. However, the permissions on the source can be retained if the file transfer is done by using for example the **rsync** command with the **-p** option (in this case, the UIDs must first be known to the Active Circle directory through synchronization with the external directory containing the UIDs).

The Unix-like permissions can be modified (using for example the **chmod** command) by connecting through a terminal session on an Active Circle share mounted using the NFS protocol.

For general background information on file system permissions, you may want to read the [Wikipedia article](#).

Access Control Lists (ACL)

The third and final verification is done in the ACLs of the files and folders. This includes all the NTFS file permissions from Windows as well as the specific [Active Circle permissions](#) (*View* and *History*).

By default, the ACLs are defined as follows:

- The ACLs at the root of the share grant full access to everyone.
- The ACLs of new files and folders are inherited from their parent folder.

For example, a file can have Unix permissions granting full access to everyone (rwxrwxrwx), while the file's ACLs only grant access to a single user. In this case, the file will only be accessible to the single user defined in the ACLs. In the same way, a user will not be able to read a file whose ACLs grant full control to everyone, if the Unix permissions don't allow such access.

It is possible to transfer files and folders while preserving their existing ACLs. For this purpose you will need to use a tool that provides these capabilities, such as RoboCopy or Xcopy. In addition, certain prerequisites must be respected. For more information, see [Appendix E, Migrating Files to Shares](#).

The file and folder permissions defined by the ACLs can be modified using different methods. The permissions specific to Active Circle are *View* and *History*, while the Windows ACLs consist of all the other file permissions. The table below shows which methods can be used for each permission category.

Table 18.3. Editing ACLs

Editing method	Permission category	
	Windows	Active Circle
Windows Explorer or Resource Kit commands on a mapped network drive	✓	--
Active Circle File Explorer	✓	✓
Active Circle Command Line (acacl command)	✓	✓
Active Circle API	✓	✓

For more information on editing Windows ACLs, see [Appendix D, Managing Windows ACLs](#).

Active Circle Permissions

There are two permissions that are specific to Active Circle; **History** and **View**. Each of these allow an additional level of rights definitions compared to the Windows permissions.

View This permission can be used to hide files or folders from users. If the View permission is removed from a file or folder, the user will not be able to see the object in a directory listing on the share. This will override any Windows permissions, meaning that the file will be invisible to the user even if all rights are granted to the user in the Windows permissions.

History This permission can be used to prevent restoration of a previous version of a file or folder. If the History permission is removed, the **Restore** action will not work (the user will receive an "access denied" message), preventing a previous version of an object from replacing the current version.



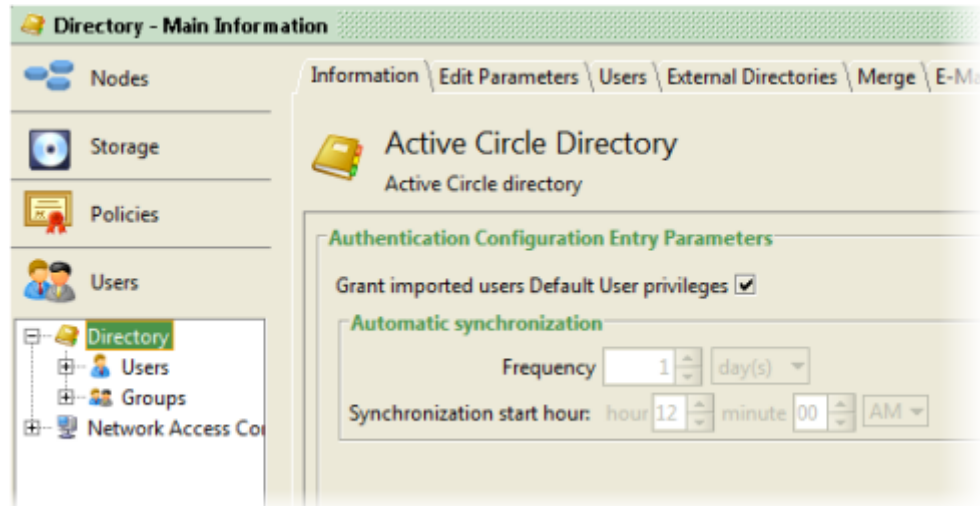
Tip

The file and folder restore functionality is available in the Active Circle File Explorer application and through the **acrestore** command on the Active Circle command line interface. For more information, please see the *Active Circle File Explorer Guide* or the *Active Circle Command Line Guide*.

Chapter 19. Directory

The **Directory** item is the first element of the **Users** view. From here, you can configure parameters affecting all users and groups in the system.

Figure 19.1. Directory Item



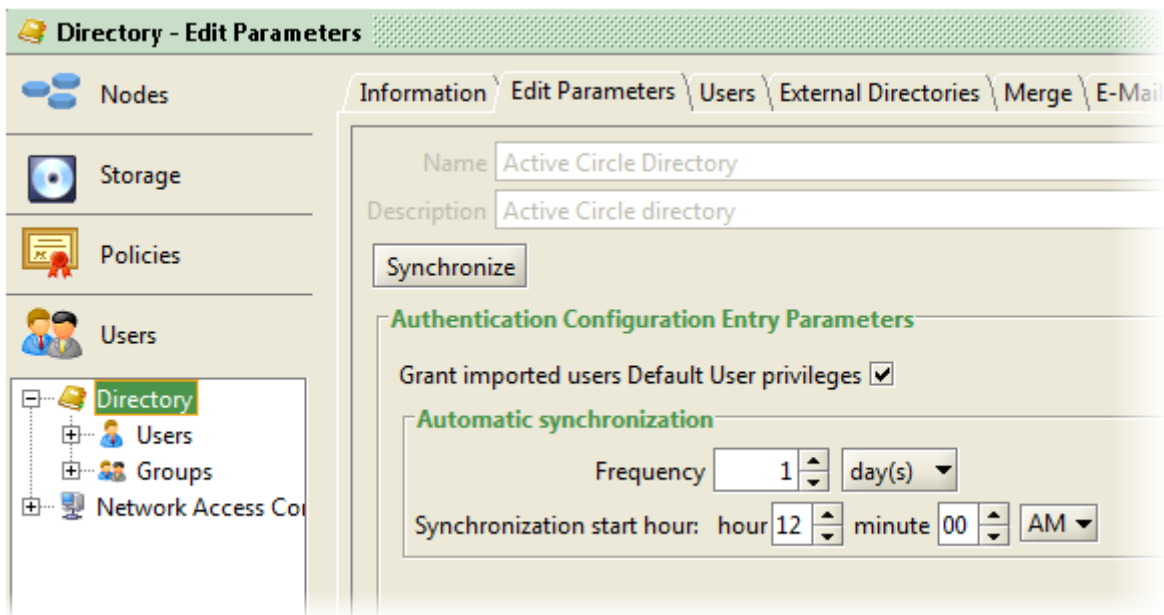
Using the settings tabs for the Directory item, you can manage the following features:

- [Synchronization](#)
- [Users](#)
- [External Directories](#)
- [Merging Users and Groups](#)
- [E-mail Addresses for Notifications](#)

Synchronization

The **Edit Parameters** tab of the **Directory** item allows you to modify parameters that impact Active Circle Directory synchronization.

Figure 19.2. Edit Parameters tab



Use the **Synchronize** button to run a manual synchronization of the directories.



Note

Synchronization is *uni-directional*: the Active Circle directories are updated from the external sources; Active Circle does not update the external directories.

The option **Grant imported users Default User privileges** automatically grants users imported from external directory services the privileges associated with the **Default User**. For more information, please see [“Users”](#).

Automatic Synchronization

The Directory is synchronized each time the Active Circle service is started on the node. You can configure the system to synchronize automatically with the external directory services on a regular basis. The default setting is to synchronize every day at midnight. You may change this by modifying the following parameters:

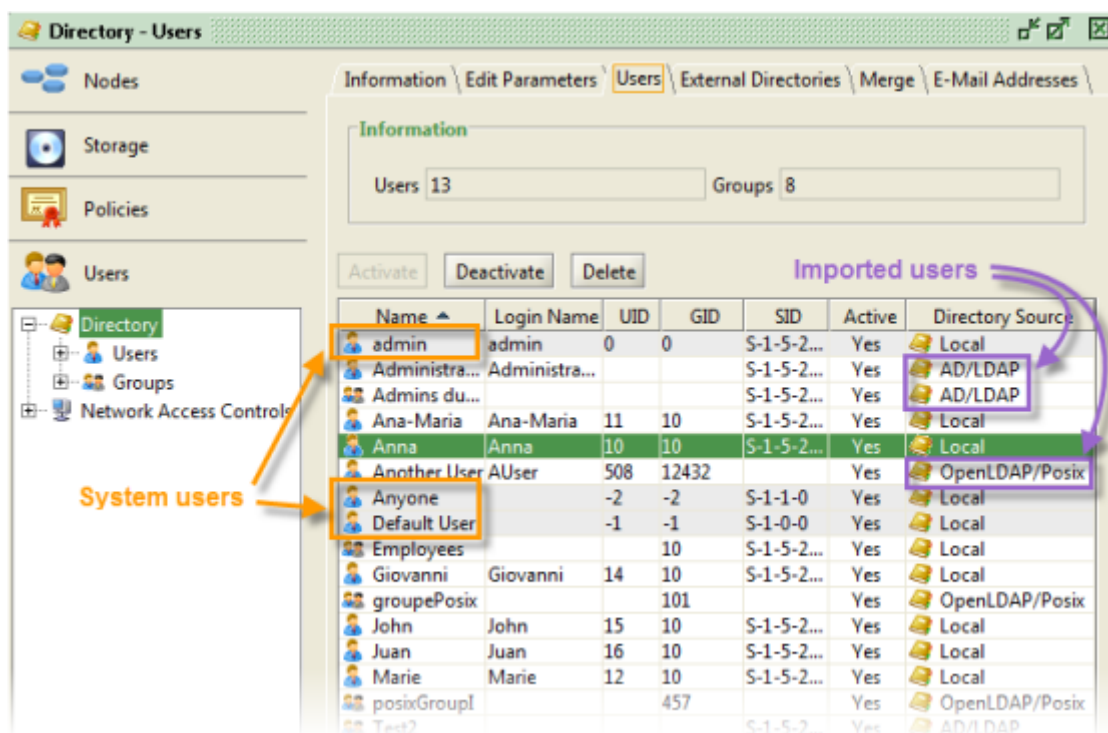
- **Frequency** — The number of days between two automatic synchronizations.
- **Synchronization start hour** — The precise **Hour** and **Minute** of the day at which the synchronization should be run.

Click the **Apply** button to validate your changes; click **Cancel** to return them to their original values.

Users

The **Users** tab presents a list of all the users and groups in the system. Three users — *admin*, *Anyone*, and *Default User* — are system users that cannot be deleted or deactivated.

Figure 19.3. Users tab



At the top of the tab are two fields providing information about the total number of users and groups in the system.

There are three buttons for managing user access to the system:

Activate

Activate the selected user(s) and/or group(s). Note that users are activated by default when created.

Deactivate

Deactivate the selected user(s) and/or group(s).

Delete

Delete the selected user(s) and/or group(s).



Note

The three users `admin`, `Anyone`, and `Default User` are system users and cannot be modified.

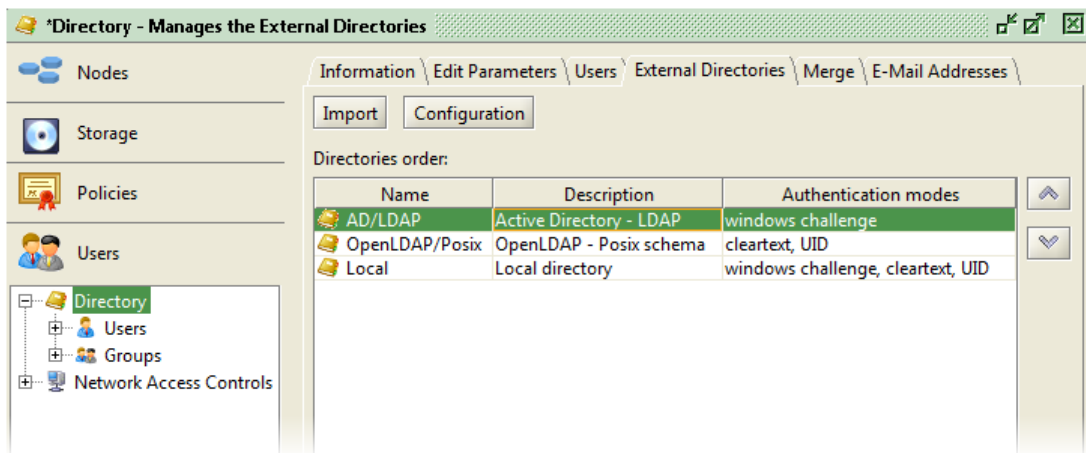
For information about how to create and modify users and groups, see [Chapter 21, Users and Groups](#).

External Directories

In order to facilitate the task of managing users and groups, Active Circle provides a system for importing from and synchronizing with external user directories.

Use the **External Directories** tab to import, configure and specify the order of precedence for the directories. You can access these settings by clicking on the **Directory** element in the **Users** view.

Figure 19.4. External Directories tab



Directory types and limitations

There are three categories of directories that can be imported:

- **LDAP (Active Directory, OpenLDAP, SunONE, etc.)**
- **NT4 / Samba (Domain or Server)**
- **NIS**

Active Circle currently supports the import of a total of *two directories*. If two directories have already been imported, you will not be able to start the wizard to import another one (you will receive a warning message instead). There are certain limitations on the combinations of external directory types that can co-exist in Active Circle:

- Only one directory based on the Windows challenge mode of authentication is allowed (for example, only one Active Directory). If special circumstances require the lifting of this limitation on your system, please contact Active Circle Support.
- Only one directory based on the UID mode of authentication is allowed (for example, only one NIS directory).
- *LDAP directories:*
 - Only the simple authentication mode is allowed if a Windows challenge compatible directory already exists.
 - In simple mode, if a UID compatible directory has already been imported, only the profiles "Active Directory" and "Simple profile using inetOrgPerson class" are allowed.

Detailed examples of the different import processes are provided in [Chapter 20, Importing External Directories](#).



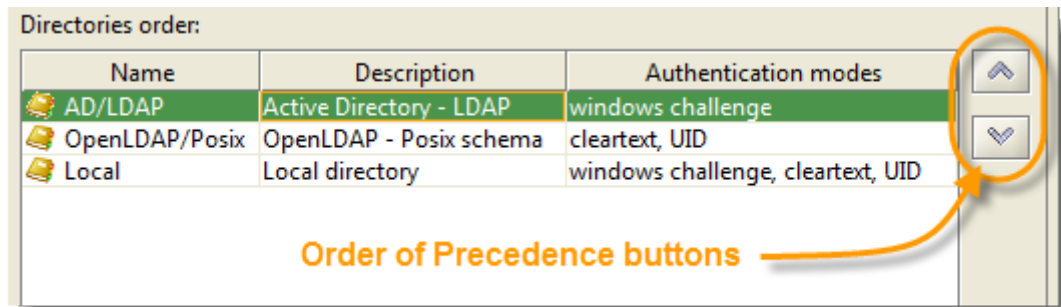
Note

Functionality for deleting a directory is by default not activated. This is because the deletion of a directory may affect the security and access rights of your data. Before considering such a step, you should have a good knowledge of the access control entries of your files which may refer to users in the directory. If you need to delete a directory, please contact Active Circle Support.

Order of Precedence

Once you have imported an external directory, the order in which they appear in the list under the **External Directories** tab plays an important role in the way in which the user accounts are verified. Use the arrow buttons to the right of the list of directories to change the order in which the directories are searched during authentication requests:

Figure 19.5. Directory Order of Precedence



- **CIFS / Windows challenge:** During authentication requests from CIFS user access, only the first Windows challenge compatible directory on the list will be searched. If the user is found in the directory at the top of the list, the user is authenticated. If an external directory using Windows challenge (Active Directory or OpenLDAP with Samba) is at the top (has precedence), any *Local* users who do not belong to the external directory will *not* be authenticated when trying to access a share through Windows (CIFS). In this case the user could be merged with a user from the external directory to allow authentication. Conversely, if the local directory has precedence, only *merged* users from the external directory will be authenticated.
- **NFS / UID:** In the case of user access through the NFS protocol, UID authentication will be used. If an external UID-compatible directory exists in addition to the local one, both directories will be searched, regardless of the order. If the user is not found in the first directory, the system continues searching in the next one.
- **FTP:** When a user tries to access a share through FTP, the system will attempt to authenticate using NTLM. Only the first NTLM-compatible directory on the list will be searched. If the user is not found, Active Circle will instead search for the user in any plaintext-compatible directories on the list. Both the local directory and Active Directory are NTLM-compatible.

The Active Circle Directory ("**Local**") supports all the above mentioned authentication methods.

Configuration

The **Configuration** button allows you to modify the parameters for synchronizing Active Circle with the external directory. The parameters here are the same as those defined during the directory import process. The fields which are displayed, vary depending on the type of directory. Some sample configuration screens are shown below:

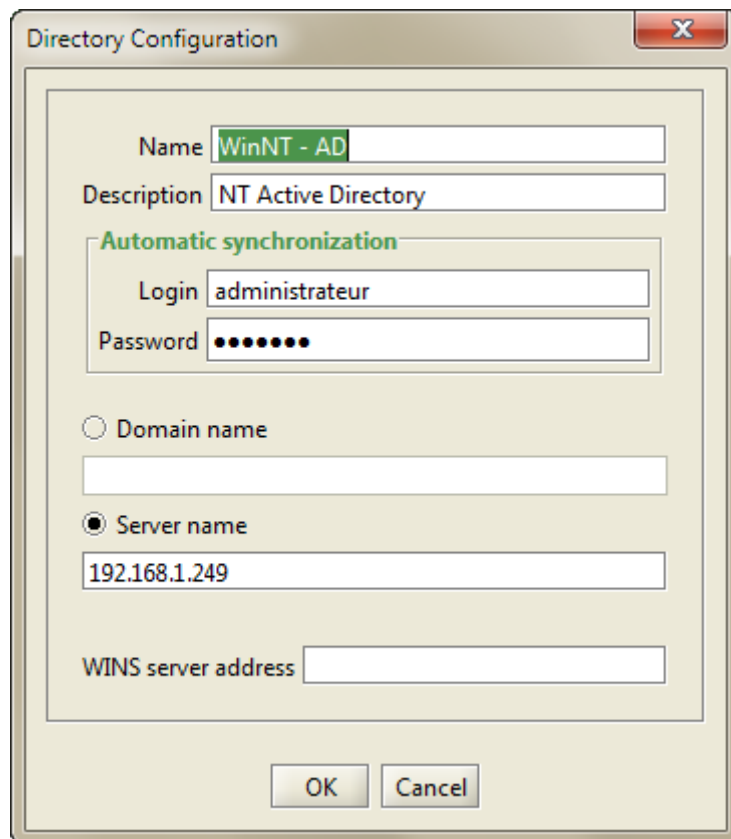
Figure 19.6. Configuration – OpenLDAP with Posix Schema (with Proxy Agent, no SASL)

The screenshot shows a 'Directory Configuration' window with a title bar and a close button. The window is divided into several sections:

- Name:** OpenLDAP/Posix
- Description:** OpenLDAP - Posix schema
- Automatic synchronization:**
 - Login:** user
 - Password:** (masked with dots)
- Server configuration:**
 - Host:** 192.168.1.249
 - Port:** 63000
 - Authentication mode:** simple (dropdown menu)
 - Search time limit (in s):** 30
 - Users search path:** dc=gargamel,dc=active-circle,dc=com
 - Groups search Path:** dc=gargamel,dc=active-circle,dc=com
 - Below these paths are buttons: Add, Modify, Delete (for both Users and Groups).
- Proxy agent configuration:**
 - ☒ Use a proxy agent for research
 - DN:** uid=ttotoo,ou=people,dc=gargamel,dc=active-circle,dc=com
 - Password:** (masked with a dot)

At the bottom right of the main configuration area is an 'Advanced >>' button. At the very bottom of the window are 'OK' and 'Cancel' buttons.

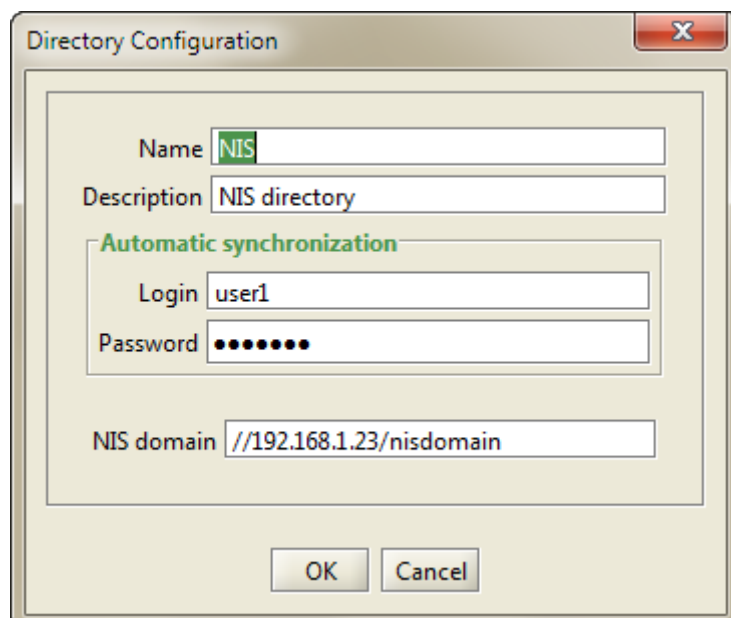
Figure 19.7. Configuration – Windows NT Mode



The 'Directory Configuration' dialog box for Windows NT Mode contains the following fields and options:

- Name:** WinNT - AD
- Description:** NT Active Directory
- Automatic synchronization:** (checked)
 - Login:** administrateur
 - Password:** (masked with dots)
- Domain name:** (radio button, unchecked)
- Server name:** (radio button, checked)
 - Server name:** 192.168.1.249
- WINS server address:** (empty field)
- Buttons:** OK, Cancel

Figure 19.8. Configuration – NIS Directory



The 'Directory Configuration' dialog box for NIS Directory contains the following fields and options:

- Name:** NIS
- Description:** NIS directory
- Automatic synchronization:** (checked)
 - Login:** user1
 - Password:** (masked with dots)
- NIS domain:** //192.168.1.23/nisdomain
- Buttons:** OK, Cancel

Some of the directory types have additional settings. If so, there is an **Advanced** button at the bottom of the configuration dialog. Click on the **Advanced** button to access more detailed settings (use the scroll bars to see all the available fields).

Figure 19.9. Advanced Configuration – LDAP with Active Directory

Directory Configuration

Users search configuration

Search filter: (&(objectclass=user)(!(objectclass=computer)))

Common name attribute: cn

Login name attribute: sAMAccountName

UID attribute:

GID attribute:

SID attribute: objectSid

Email attribute: mail

Primary Group Attribute: PrimaryGroupID

Groups search configuration

Search filter: (objectclass=group)

Name attribute: cn

GID attribute:

SID attribute: objectSid

Members attribute: member

Group membership attribute type: DN (distinguished name)

☒ Recursive groups

☒ Ignore groups of type "distribution"

OK Cancel

Merging Users and Groups

The **Merge** tab allows you to manage how user and group accounts coming from two different directory sources are handled. Specifically, it allows you to merge or undo the merge of user accounts that refer to the same individual or group.

Figure 19.10. Merge tab

Directory - Manages Merge of External Directories Users

Information | Edit Parameters | Users | External Directories | **Merge** | E-Mail Addresses

Last update: August 22, 2011 5:19:29 PM

Information

Users: 4 Groups: 1

Merge Users Merge Groups Undo Merge Refresh

Name	Local	NIS	AD / LDAP	Merge Type
Administrateur	Administrateur		Administrateur	Automatic merge
Mr User1	User1	user1		Automatic merge
test1	test1		test1	Automatic merge
Test2		Test2	Test2	Automatic merge
UserUID501		UserUL...	UserUID501	Automatic merge



Note

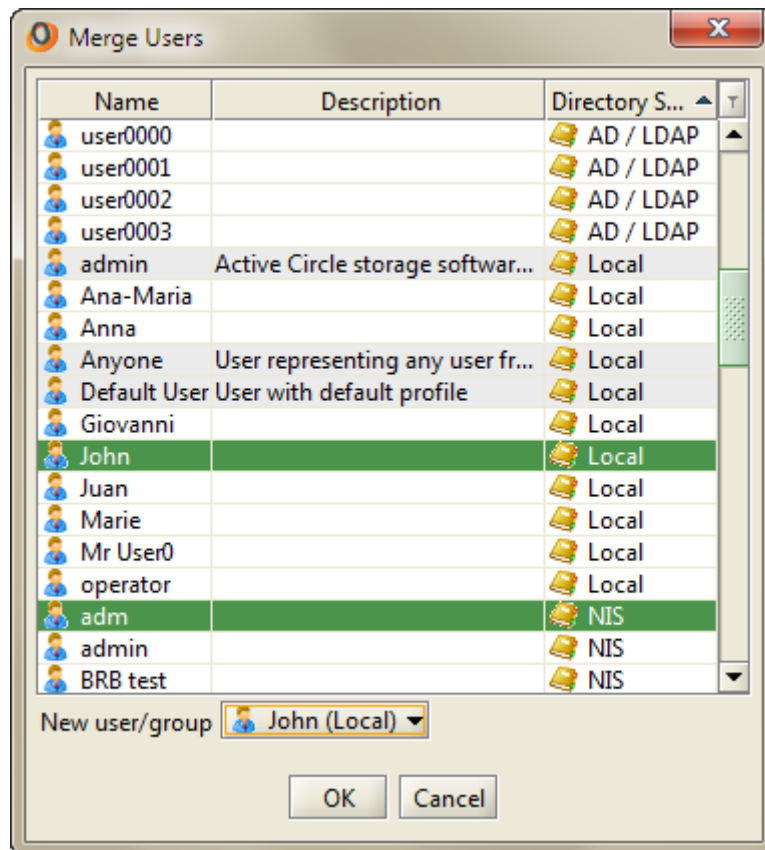
Automatic merge: If you import a directory and the system detects that for one of the users being imported there is already another user that uses the same login name, then the two user accounts will be merged into one in the Active Circle Directory. These two user accounts will then be seen as the same user.

If the automatically merged users are two separate users with identical login names, you can *undo the merge*. Please see [“Undo Merge”](#) for more information.

Merge Users

To perform a manual merge, click the **Merge Users** button. A dialog box opens displaying a list of all the users and groups in the system.

Figure 19.11. Merge Users dialog box



Note

You cannot merge users from the same external directory source.

Select the users to be merged from the list using **Ctrl+click** to select multiple entries.

Next, you must choose which of the selected users will be the primary one using the **New user/group** drop-down list. The references to the two user names that were merged are removed from Active Circle (though not from the source external directories) and this primary user will become the name of the user in the **Local** directory.

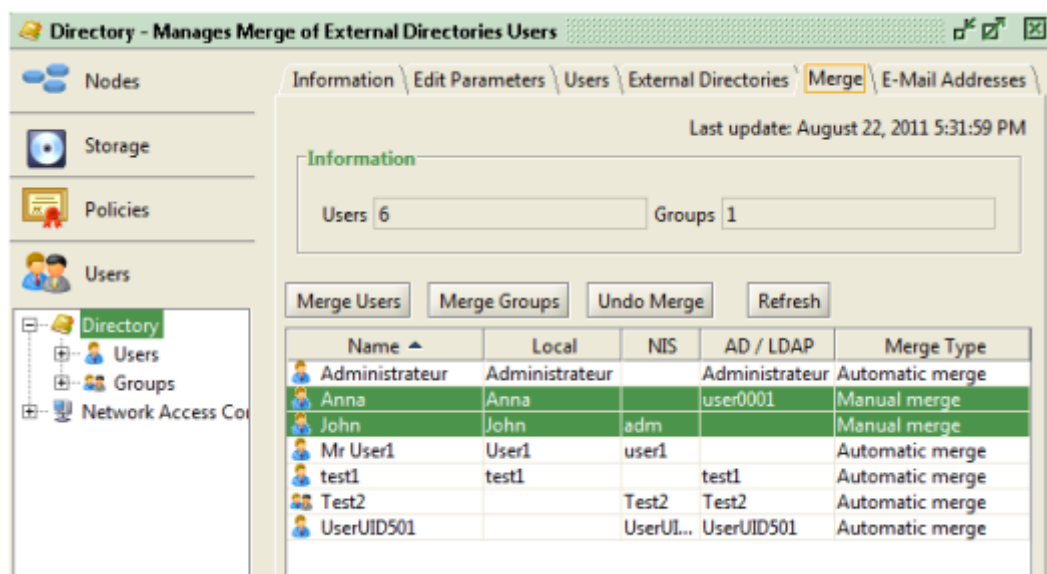


Important

The choice of primary user is also significant if you want to undo a merge of a user. When you unmerge, the two original users are recreated in the lists for the external directories and the primary user will then become the owner of any objects created by the two users while they were merged. For more information, please see [“Undo Merge”](#).

When you are finished, click **OK** to merge the users and close the dialog box. You will then see the merged user in the list. If it does not immediately appear, click the **Refresh** button.

Figure 19.12. Merge tab with manually merged users



Merge Groups

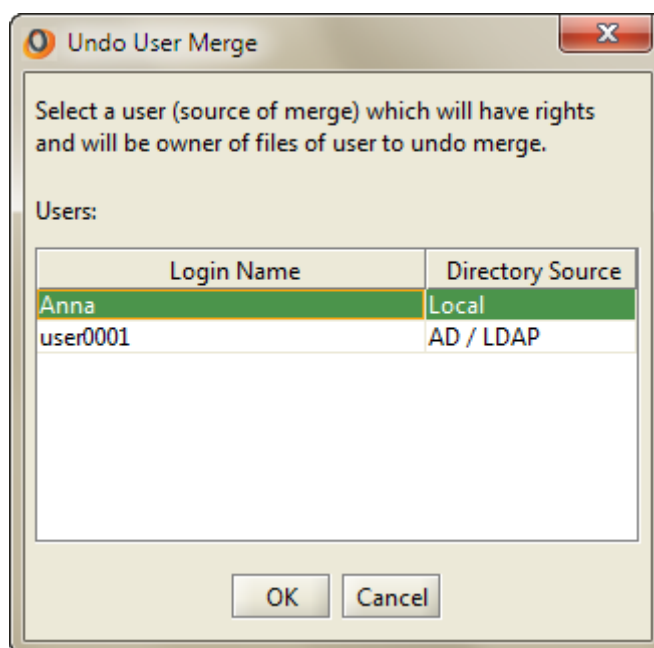
The **Merge Groups** button allows you to merge groups. It functions exactly like the **Merge Users** button described above.

Undo Merge

You can use the **Undo Merge** button to undo both automatic and manual merges. To undo a merge, select a user and click the button **Undo Merge**.

A dialog box opens listing the users who were the source of the merge. You must choose one of these accounts to be the sole owner of any objects created in the system during the time the accounts were merged.

Figure 19.13. Undo User Merge dialog box



The user you selected as sole owner will keep all the user profiles of the previously merged user. The user you didn't select as owner will be left only with the user profile called "User".



Important

When undoing an automatic merge, remember that the two user accounts coming from two separate directories had identical login names. Because Active Circle will not allow two users to share a login name, the undo merge process will automatically rename the login of the user *not selected* above. That is, the user who will not be the owner. The name of the source directory will be appended to the login name.

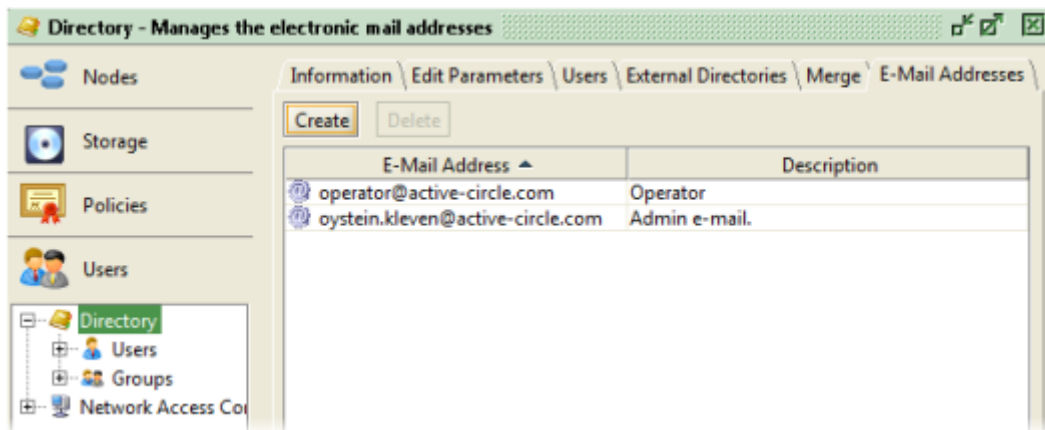
Refresh

The **Refresh** button updates the list of merged users and groups after performing a merge or undoing a merge.

E-Mail Addresses


The **E-Mail Addresses** tab provides you with a way to simplify system event notification by allowing you to create a list of e-mail addresses. These addresses can then be registered to receive automated Supervision Notes notifying users and/or groups of system events.

Figure 19.14. E-Mail Addresses tab



Click the Create button, enter an e-mail address and optionally a description, and then click OK.

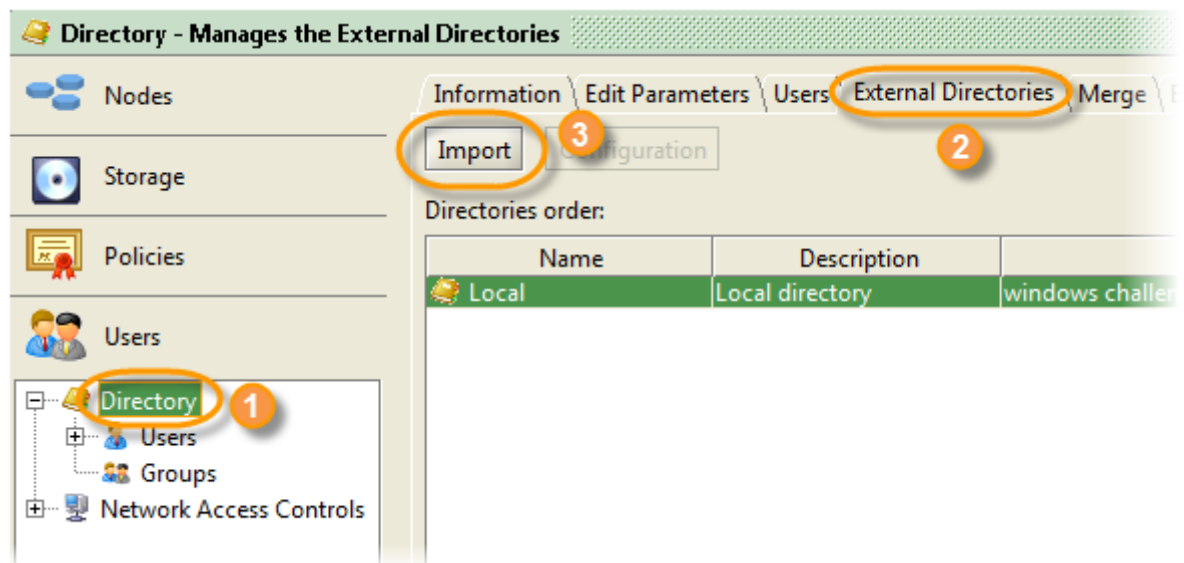
Figure 19.15. Create an E-Mail Address



Chapter 20. Importing External Directories

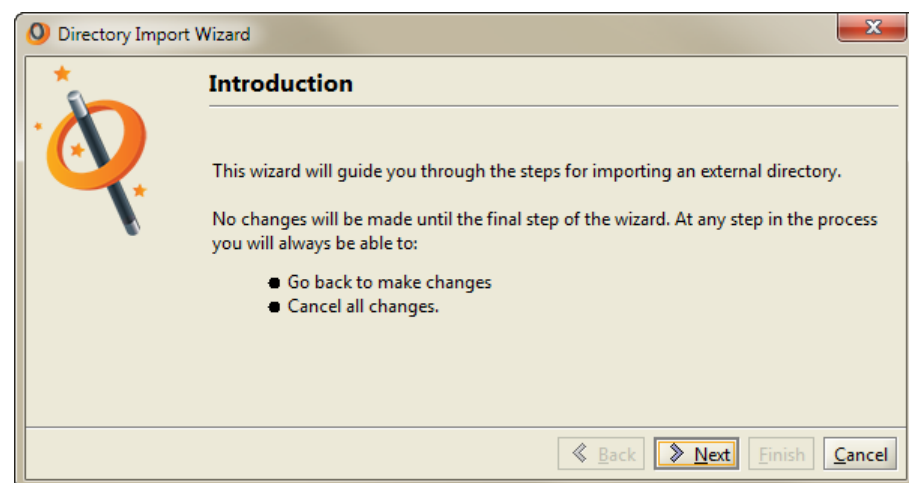
This chapter provides step-by-step instructions for importing the different types of supported external directories. You can import a total of two external directories in Active Circle.

Figure 20.1. Import External Directories



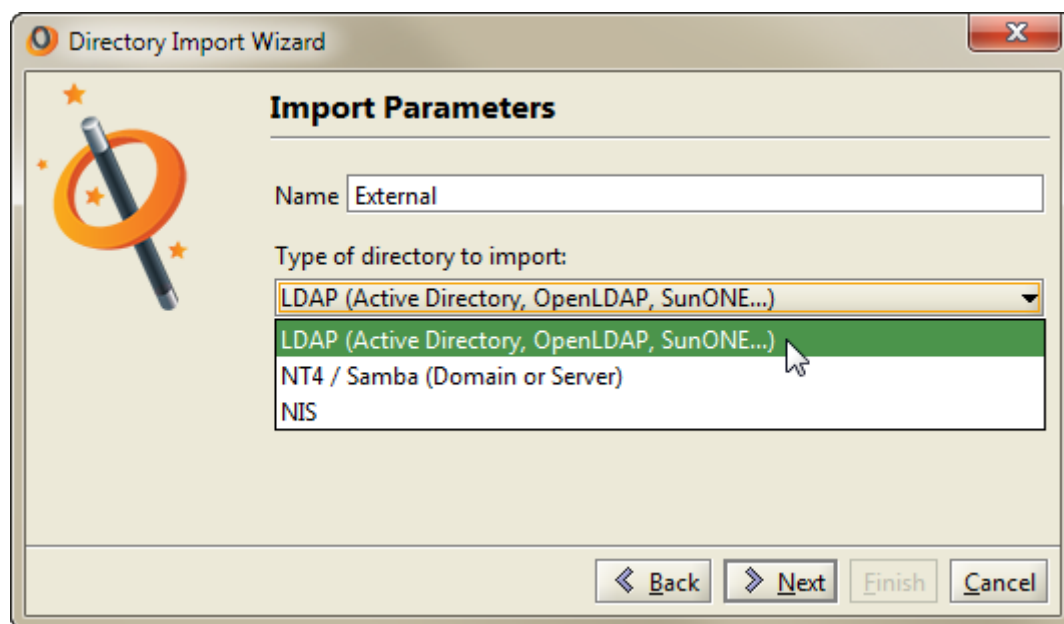
To launch the **Directory Import Wizard**, select the **Directory** element in the **Users** view, click on the **External Directories** tab, and then click the **Import** button. The wizard introduction screen appears.

Figure 20.2. Introduction screen



Click **Next** to proceed to the **Import Parameters** screen where you will choose the type of directory to import.

Figure 20.3. Import Parameters



Provide a **Name** that will be used to identify the directory to be imported within Active Circle, and select the directory type from the drop-down menu **Type of directory to import**. The options are:

- **LDAP (Active Directory, OpenLDAP, SunONE...)**
- **NT 4 / Samba (Domain or Server)**
- **NIS**

Select the directory type appropriate to your external directory and then click the **Next** button to continue. Your choice here will determine the path of the remaining steps of the **Directory Import Wizard**. Example import procedures corresponding to the various choices are given in the sections that follow. They are:

- [OpenLDAP Directory with POSIX Schema](#)
- [Active Directory Imported in LDAP Mode](#)
- [Active Directory Imported in Windows NT Mode](#)
- [NIS Directory](#)



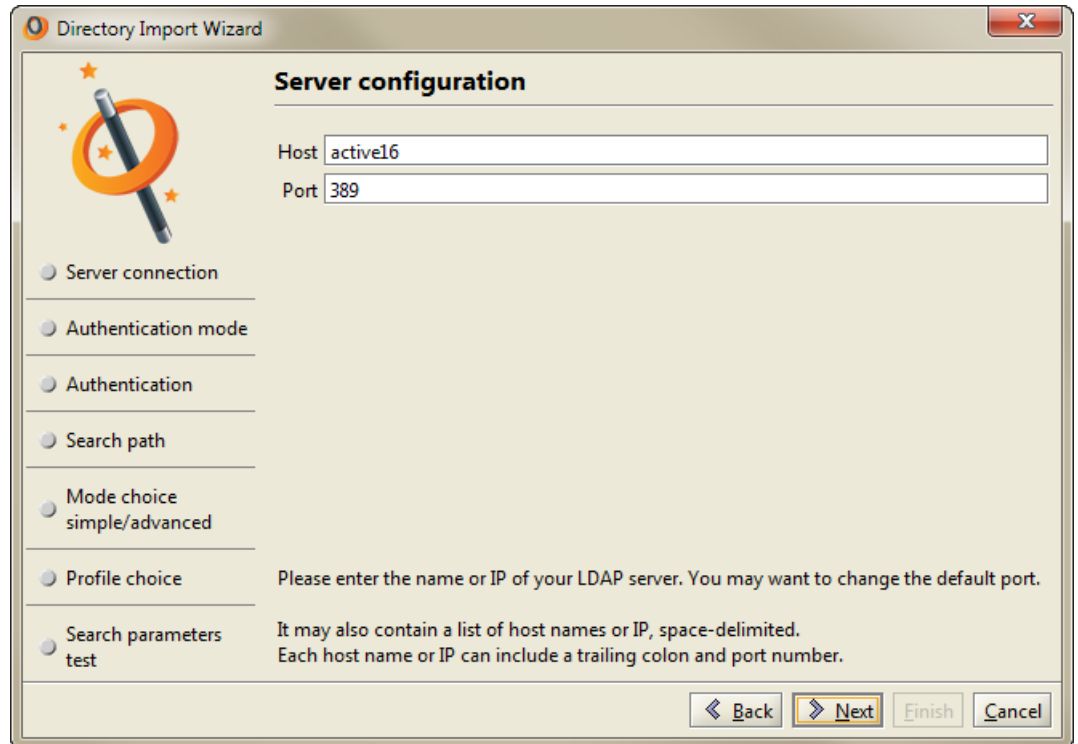
Note

The types of directories available for import depend on whether another directory of a certain type has already been imported. For details on which types of external directories can co-exist in Active Circle, see [“Directory types and limitations”](#).

OpenLDAP Directory with POSIX Schema

This section describes the procedure for importing an OpenLDAP directory that conforms to the POSIX schema.

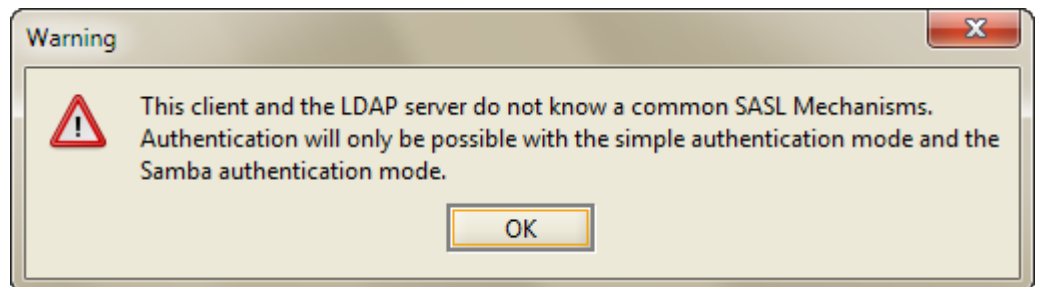
1. Server Configuration



The screenshot shows the 'Directory Import Wizard' window with the 'Server configuration' tab selected. On the left, a vertical list of steps includes 'Server connection', 'Authentication mode', 'Authentication', 'Search path', 'Mode choice simple/advanced', 'Profile choice', and 'Search parameters test'. The 'Server connection' step is currently active. The main area contains two input fields: 'Host' with the value 'active16' and 'Port' with the value '389'. Below these fields, there is explanatory text: 'Please enter the name or IP of your LDAP server. You may want to change the default port. It may also contain a list of host names or IP, space-delimited. Each host name or IP can include a trailing colon and port number.' At the bottom right, there are four buttons: 'Back', 'Next' (highlighted with a yellow border), 'Finish', and 'Cancel'.


In this screen, type the hostname or IP address of your LDAP server and its port number. Click the **Next** button to continue.

If Active Circle detects that it does not share an SASL mechanism with the directory, a warning message is displayed:



The message informs you of the available authentication modes to choose from in the next step of the procedure.

2. Authentication



Select an **Authentication mode**. The options are:

SASL

Simple Authentication and Security Layer.

Note that this option is only available if the LDAP server and the Active Circle client share an SASL authentication mechanism.

samba

Authentication mode for an OpenLDAP directory in a Windows/Samba environment.

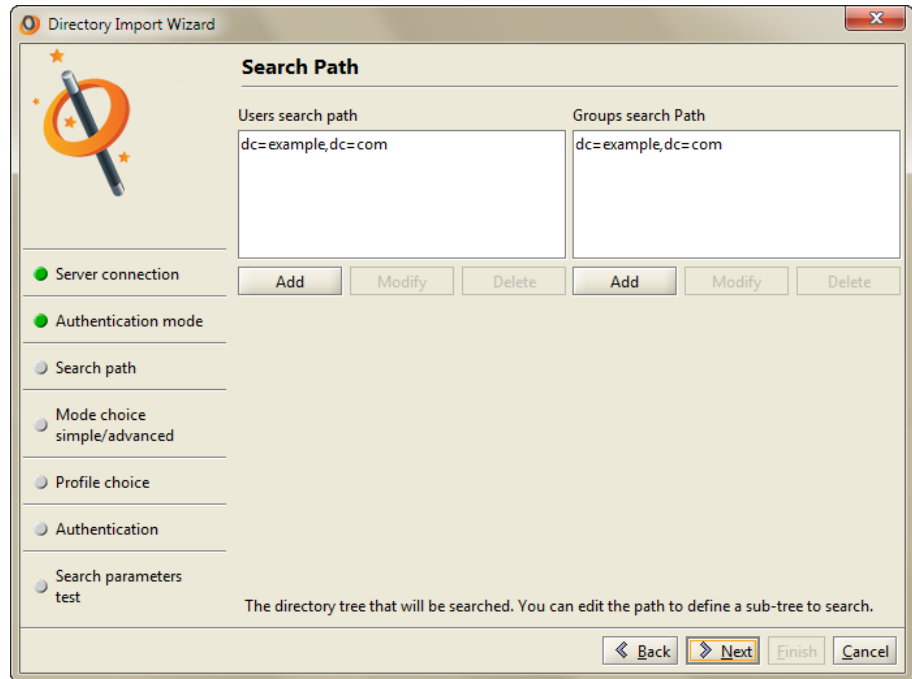
This mode is not compatible with a POSIX schema.

simple

This is the simplest form of authentication for LDAP directories. The login credentials are passed over the network in plaintext. This mode does not support NTLM (IWA).

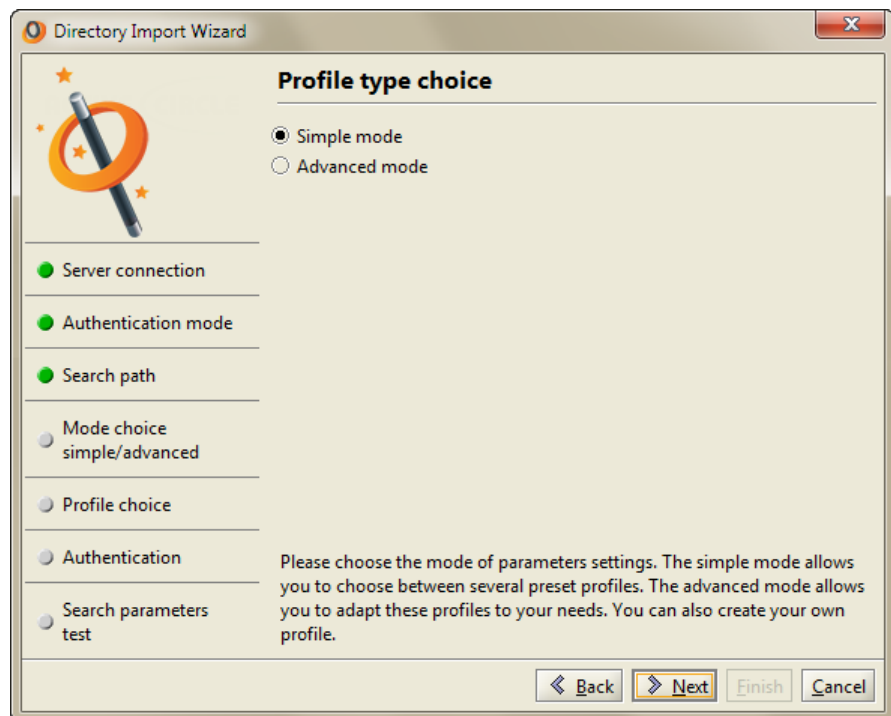
In either `samba` or `simple` mode, you can use a proxy agent to search the directory by providing the Distinguished Name (**DN**) and password. If no proxy is defined, the searches are made anonymously.

3. Search Path



This screen displays the roots of the directory search paths: the **Users search** path and the **Groups search** path. You can either accept the default or modify the search paths using the Add, Modify, and Delete buttons.

4. Profile Type Choice



In this screen you must choose how to import the user profiles. There are two options:

Chapter 20. Importing External Directories

Simple mode

Allows you to select from a set of pre-defined profiles.

Advanced mode

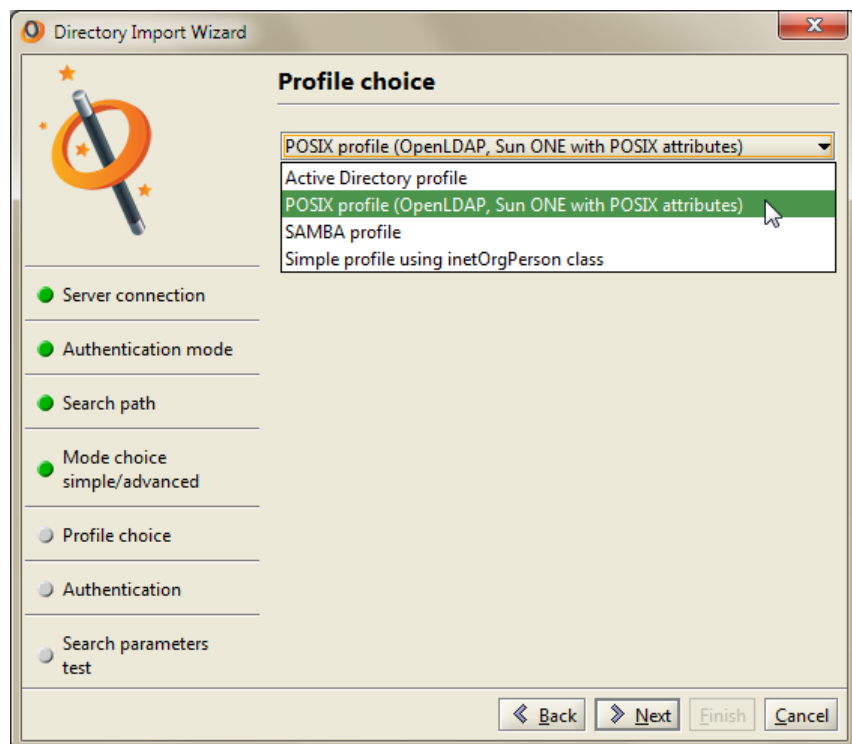
Allows you to adapt the profile import configuration to the specific needs of your organization.

The choice you make here will decide which screen will follow. Click the **Next** button to continue.

5. Profile choice

The screen presented in this step depends on your selection in the previous screen.

- **Profile choice: Simple Mode**



Select one of the predefined profile types:

- **Active Directory profile**
- **POSIX profile (OpenLDAP, SunONE with POSIX attributes).**
Not available if selected authentication mode is "samba".
- **SAMBA profile**
- **Simple profile using inetOrgPerson class**

Click **Next** to continue.

- **Profile choice: Advanced Mode**

The advanced mode gives you access to the same list of profile choices as the simple mode with one option added:

- **User defined profile**

Chapter 20. Importing External Directories

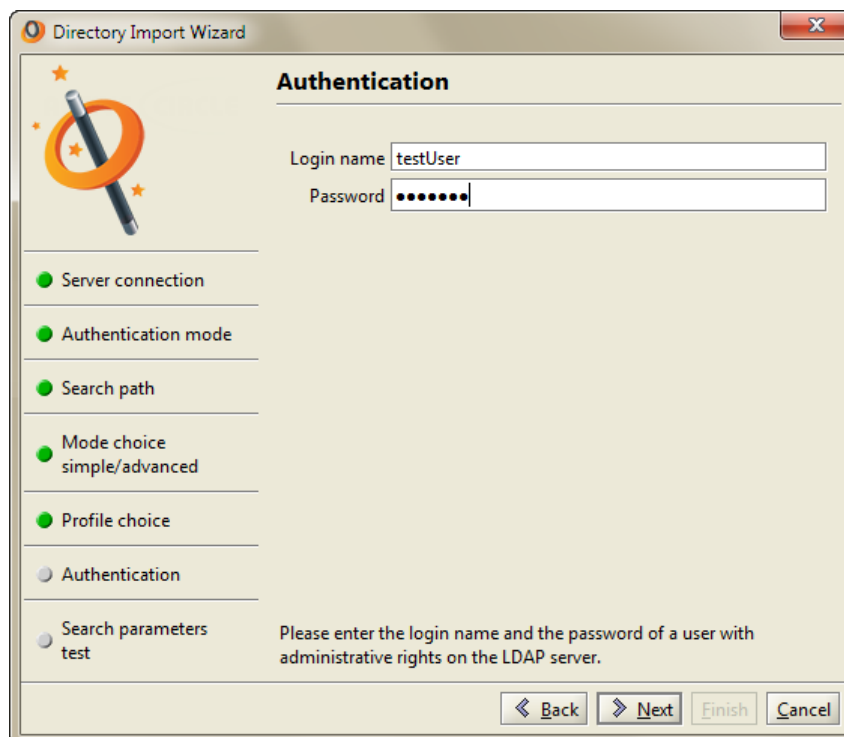
If you choose this profile, all the fields will be blank, and you must manually enter the information used for importing the external directory. If you use one of the pre-configured profiles, the predefined values will be provided in the wizard:

The screenshot shows the 'Directory Import Wizard' window, specifically the 'Profile choice' step. On the left is a sidebar with a list of steps: 'Server connection', 'Authentication mode', 'Search path', 'Mode choice simple/advanced', 'Profile choice' (which is selected and highlighted with a green dot), 'Authentication', and 'Search parameters test'. The main area is titled 'Profile choice' and contains a 'Profiles' dropdown menu set to 'POSIX profile (OpenLDAP, Sun ONE with POSIX attributes)'. Below this are three configuration sections: 'Server configuration' with fields for 'Name' (active16), 'Port' (389), and 'Search time limit (in s)' (30); 'Users search configuration' with fields for 'Search filter' ((objectclass=posixAccount)), 'Common name attribute' (cn), 'Login name attribute' (uid), 'UID attribute' (uidNumber), 'GID attribute' (gidNumber), 'SID attribute' (blank), 'Email attribute' (blank), and 'Primary Group Attribute' (blank); and 'Groups search configuration' with fields for 'Search filter' ((objectclass=posixGroup)), 'Name attribute' (cn), 'GID attribute' (gidNumber), 'SID attribute' (blank), 'Members attribute' (memberUid), and 'Group membership attribute type' (Login name attribute). There are also checkboxes for 'Recursive groups' (unchecked) and 'Ignore groups of type "distribution"' (checked). At the bottom right are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

In the advanced mode, you can modify the values suggested by the predefined profile you have chosen. If you make no modifications to the predefined profile values, the advanced mode will use the same settings as if you had selected the simple mode. When you have finished configuring the profile options, click **Next**.

6. Authentication

This screen enables you to provide the connection information of an account with access to the external directory.



Authentication

Login name: testUser

Password:

Progress indicators:

- Server connection
- Authentication mode
- Search path
- Mode choice: simple/advanced
- Profile choice
- Authentication
- Search parameters: test

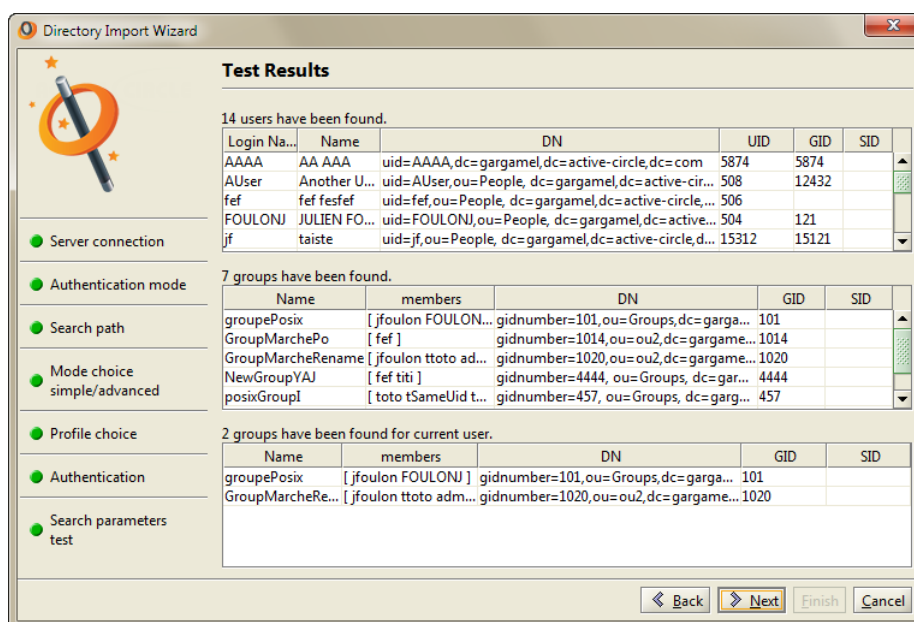
Please enter the login name and the password of a user with administrative rights on the LDAP server.

Buttons: Back, Next, Finish, Cancel

When you click **Next**, Active Circle will check the credentials. If an incorrect login has been entered, a warning message will appear.

7. Test Results

This screen provides the results of the test access to the external directory using the credentials specified in the previous screen. The tables show the lists of users and groups found in the directory, and the groups (if any) the current user belongs to.



Test Results

14 users have been found.

Login Na...	Name	DN	UID	GID	SID
AAAA	AA AAA	uid=AAAA, dc=gargamel, dc=active-circle, dc=com	5874	5874	
AUser	Another U...	uid=AUser, ou=People, dc=gargamel, dc=active-cir...	508	12432	
fef	fef fesfef	uid=fef, ou=People, dc=gargamel, dc=active-circle, ...	506		
FOULONJ	JULIEN FO...	uid=FOULONJ, ou=People, dc=gargamel, dc=active...	504	121	
jf	taiste	uid=jf, ou=People, dc=gargamel, dc=active-circle, d...	15312	15121	

7 groups have been found.

Name	members	DN	GID	SID
groupePosix	[jfoulon FOULONJ]	gidnumber=101, ou=Groups, dc=garga...	101	
GroupMarchePo	[fef]	gidnumber=1014, ou=ou2, dc=gargame...	1014	
GroupMarcheRename	[jfoulon ttoto adm...	gidnumber=1020, ou=ou2, dc=gargame...	1020	
NewGroupYAJ	[fef titi]	gidnumber=4444, ou=Groups, dc=gar...	4444	
posixGroupI	[toto tSameUid t...	gidnumber=457, ou=Groups, dc=garg...	457	

2 groups have been found for current user.

Name	members	DN	GID	SID
groupePosix	[jfoulon FOULONJ]	gidnumber=101, ou=Groups, dc=garga...	101	
GroupMarcheRe...	[jfoulon ttoto adm...	gidnumber=1020, ou=ou2, dc=gargame...	1020	

Buttons: Back, Next, Finish, Cancel

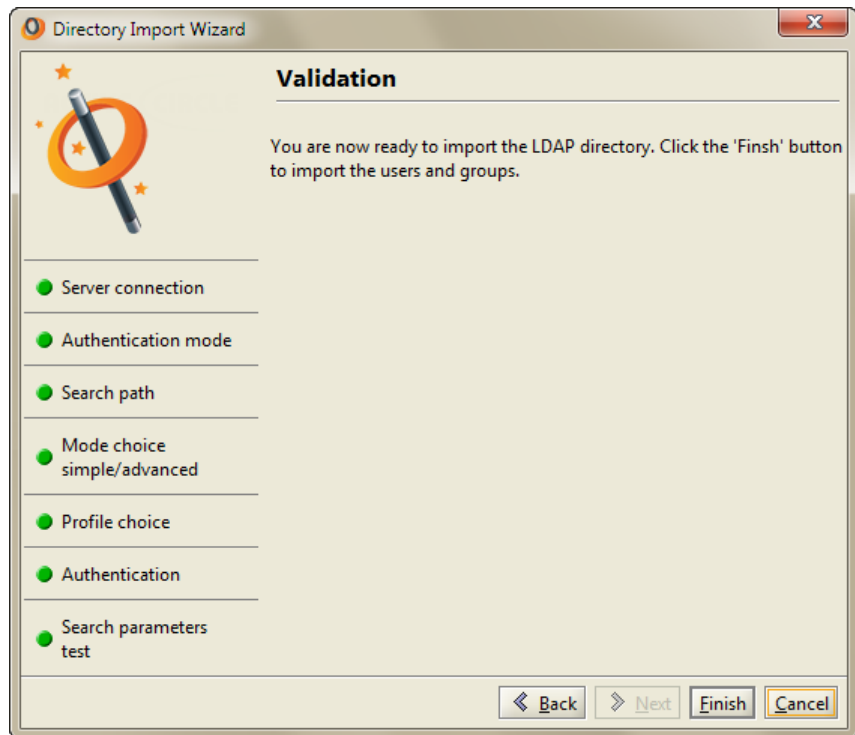


Note

The external directory has not yet been imported. Active Circle has only verified the access to that directory.

8. Validation

This is the final step in the import process. Click on **Finish** to begin the import.

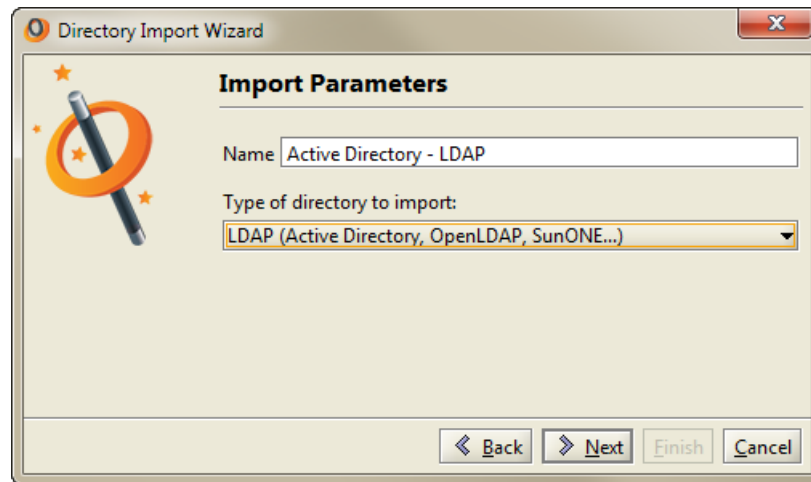


You can view or modify the connection information for the imported directory by selecting it in the **External Directories** tab and clicking the **Configuration** button.

Active Directory Imported in LDAP Mode

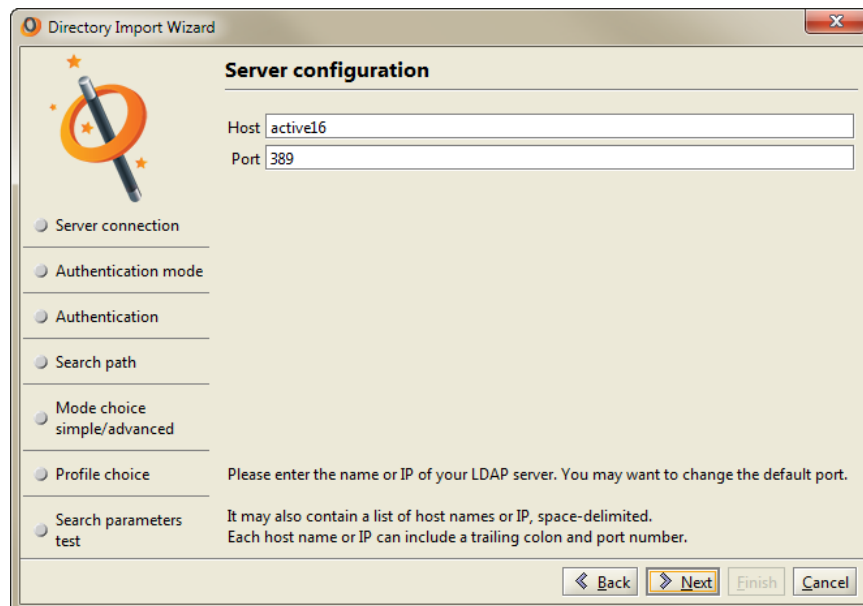
This section describes the procedure for importing an Active Directory in LDAP mode.

1. Import Parameters



Type a **Name** to be used to identify the imported directory withing Active Circle. Then select the option **LDAP (Active Directory, OpenLDAP, SunONE...)** from the drop-down list labeled **Type of directory to import**. Click **Next** to continue.

2. Server Configuration



Specify the server hosting the external directory.

Host name

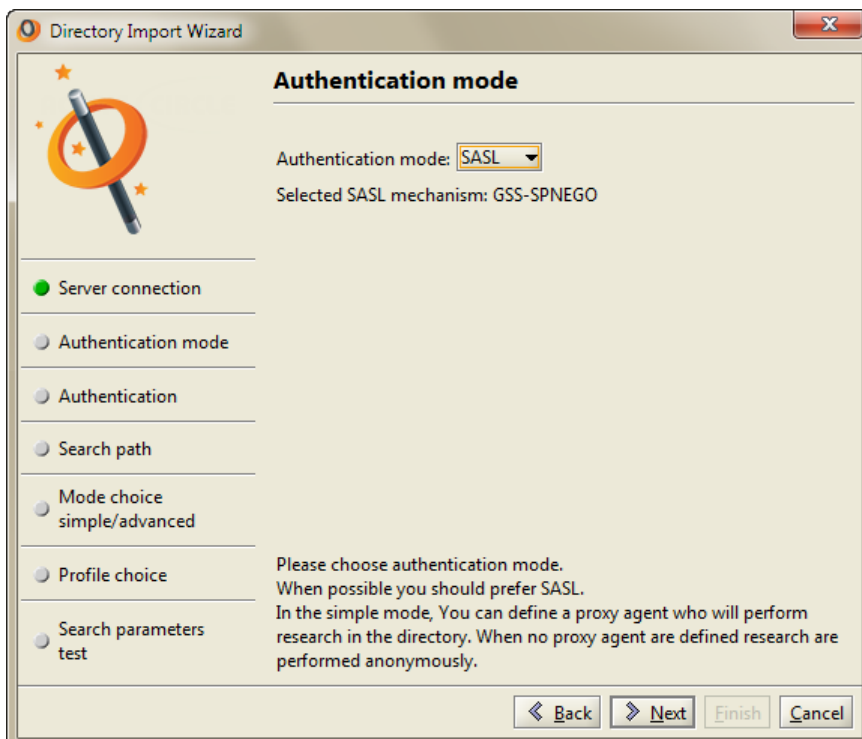
The IP address or DNS name of the server. If your environment has multiple domain controllers (primary, secondary), you can create a list here by adding the identifiers of the servers separated by a space character.

Port

The port number on which the server is listening.

Click **Next** to continue.

3. Authentication mode

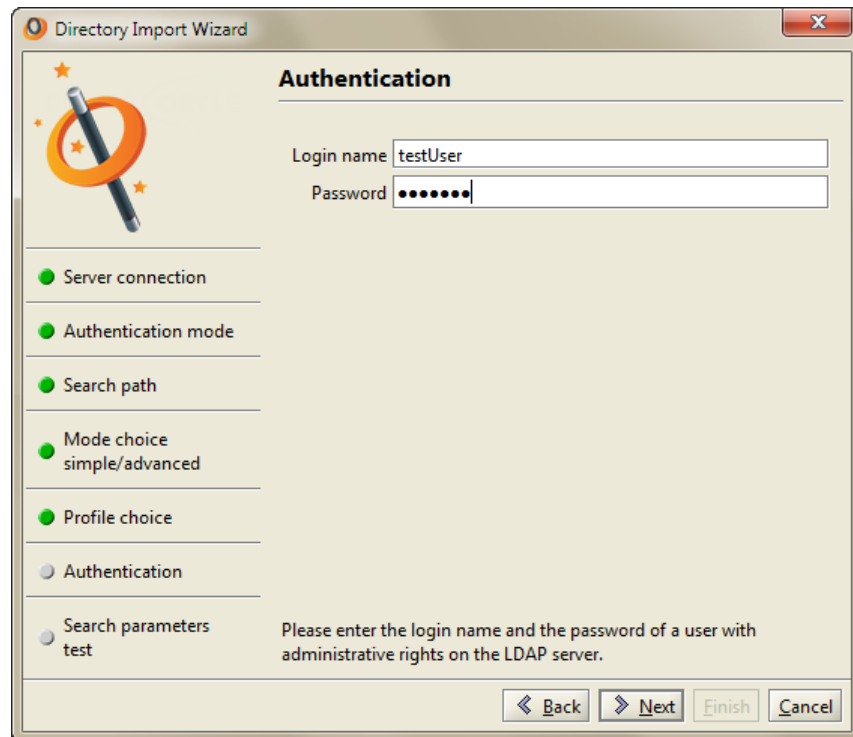


Select the authentication mechanism to be used when Active Circle connects to the external directory. There are three methods:

SASL	Simple Authentication and Security Layer. This option is only available if the LDAP server and the Active Circle client share an SASL authentication mechanism.
samba	Authentication mode for an OpenLDAP directory in a Windows/Samba environment.
simple	This is the simplest form of authentication for LDAP directories. The login credentials are passed over the network in plaintext. This mode does not support NTLM (IWA).

In our example, the LDAP server and the Active Circle client share an SASL authentication mechanism, so the **SASL** method is already selected by the wizard. Click **Next** to continue.

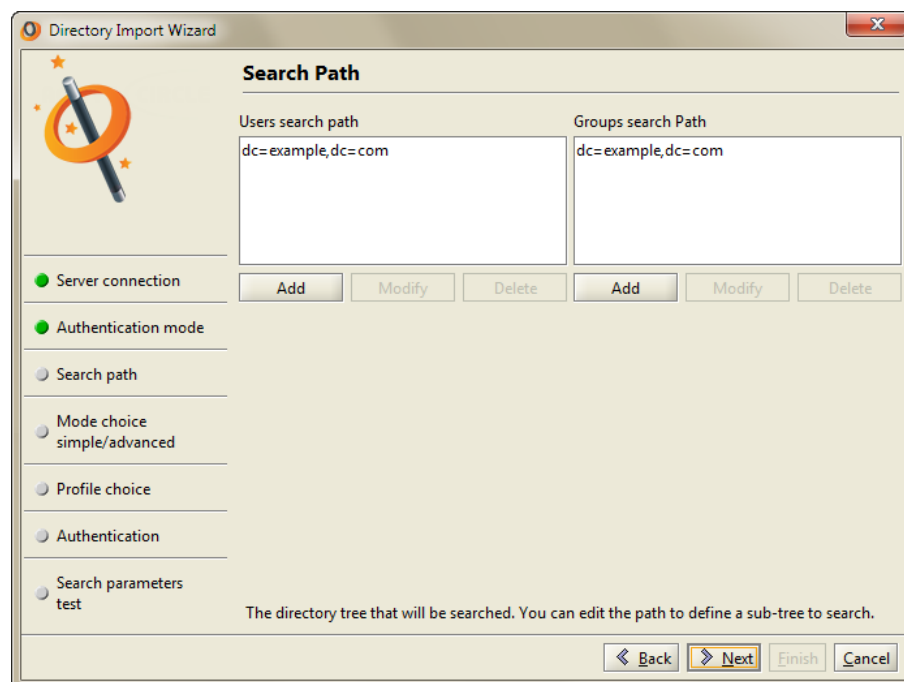
4. Authentication



The screenshot shows the 'Authentication' step of the Directory Import Wizard. The window has a title bar 'Directory Import Wizard' and a close button. On the left is a sidebar with a list of steps: 'Server connection', 'Authentication mode', 'Search path', 'Mode choice simple/advanced', 'Profile choice', 'Authentication', and 'Search parameters test'. The 'Authentication' step is currently selected. The main area is titled 'Authentication' and contains two input fields: 'Login name' with the text 'testUser' and 'Password' with masked characters '.....'. Below the input fields is a text instruction: 'Please enter the login name and the password of a user with administrative rights on the LDAP server.' At the bottom right are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

Specify a **Login name** and **Password** for an account with access to the external directory. Click **Next** to continue.

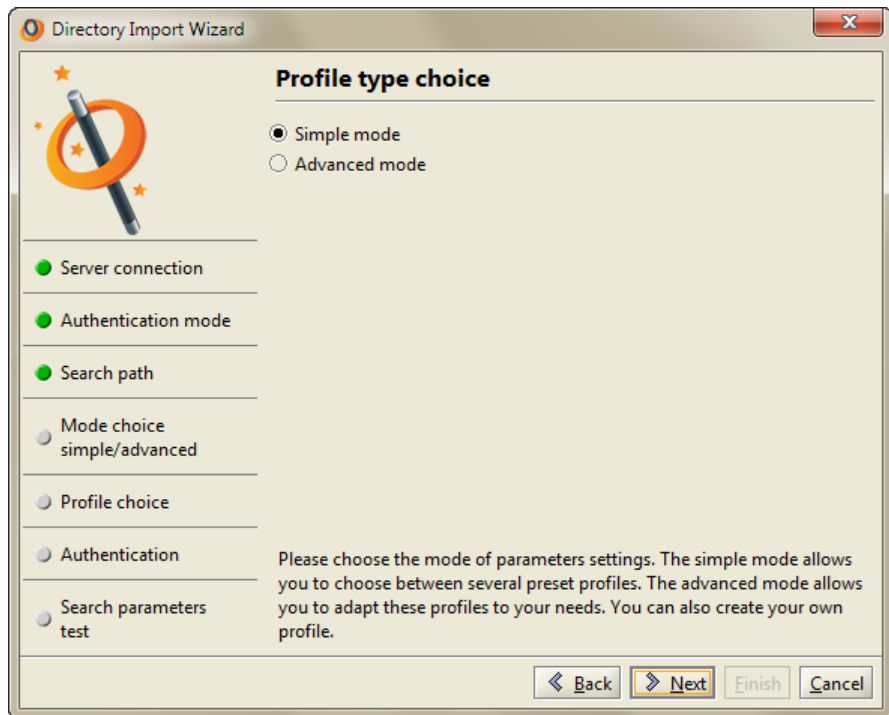
5. Search Path



The screenshot shows the 'Search Path' step of the Directory Import Wizard. The window has a title bar 'Directory Import Wizard' and a close button. On the left is a sidebar with a list of steps: 'Server connection', 'Authentication mode', 'Search path', 'Mode choice simple/advanced', 'Profile choice', 'Authentication', and 'Search parameters test'. The 'Search path' step is currently selected. The main area is titled 'Search Path' and contains two input fields: 'Users search path' and 'Groups search Path', both containing the text 'dc=example,dc=com'. Below each input field are three buttons: 'Add', 'Modify', and 'Delete'. At the bottom right are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'. A text instruction at the bottom reads: 'The directory tree that will be searched. You can edit the path to define a sub-tree to search.'

This screen displays the roots of the directory search paths: the **Users search** path and the **Groups search** path. You can either accept the default or modify the search paths using the **Add**, **Modify**, and **Delete** buttons.

6. Profile Type Choice



In this screen you must choose how to import the user profiles. There are two profile modes to choose from:

Simple mode

Allows you to select from a set of pre-defined profiles.

Advanced mode

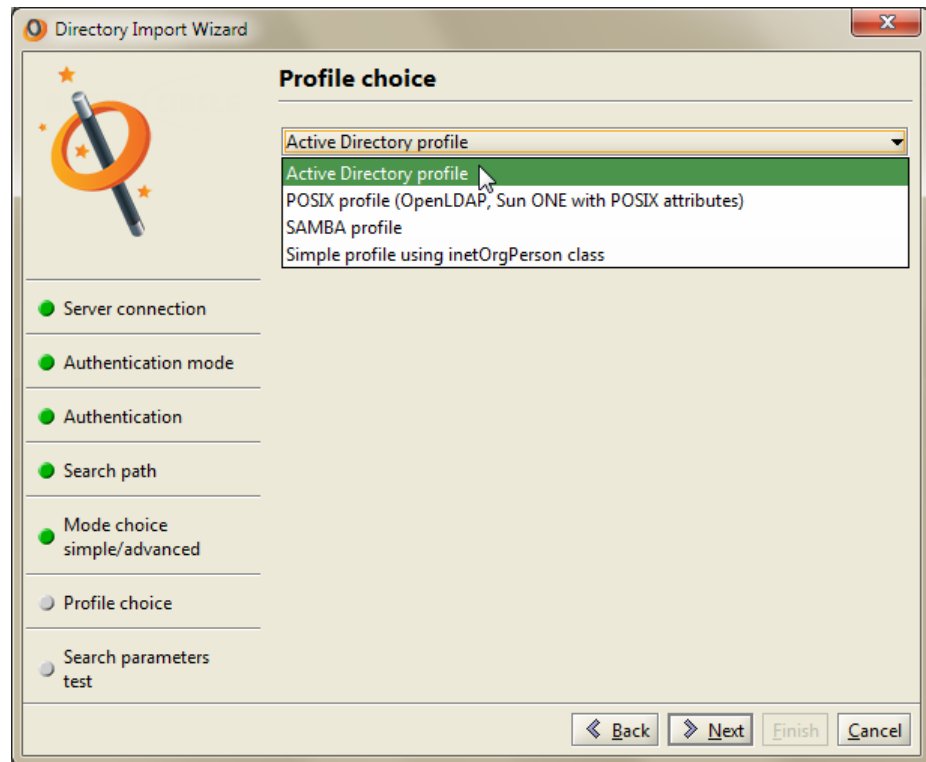
Allows you to adapt the profile import configuration to the specific needs of your organization.

The screen that follows this depends on your choice here. Click the **Next** button to continue.

7. Profile choice

The screen presented in this step depends on which profile mode you selected in the previous screen.

- **Profile choice: Simple Mode**



In this mode, you can select from one of the predefined profile types. They are:

- **Active Directory profile**
- **POSIX profile (OpenLDAP, SunONE with POSIX attributes)**
- **SAMBA profile**
- **Simple profile using inetOrgPerson class**

Select the **Active Directory profile**. Click **Next** to proceed to the **Test Results** screen.

- **Profile choice: Advanced Mode**

Directory Import Wizard

Profile choice

Profiles: **Active Directory profile**

Server configuration

Name: 192.168.2.25 Port: 389 Search time limit (in s): 30

Users search configuration

Search filter: (&(objectclass=user)!(&(objectclass=computer)))

Common name attribute: cn

Login name attribute: sAMAccountName

UID attribute:

GID attribute:

SID attribute: objectSid

Email attribute: mail

Primary Group Attribute: PrimaryGroupID

Groups search configuration

Search filter: (objectclass= group)

Name attribute: cn

GID attribute:

SID attribute: objectSid

Members attribute: member

Group membership attribute type: DN (distinguished name)

☒ Recursive groups

☒ Ignore groups of type "distribution"

Back Next Finish Cancel

The *Advanced Mode* gives you access to the same list of profile choices as the simple mode with one option added:

- **User defined profile**

If you choose this profile, all the fields will be blank, and you must manually enter the information used for importing the external directory. If you use one of the pre-configured profiles, the predefined values will be provided in the wizard.

In the advanced mode, you can modify the values suggested by the predefined profile you have chosen. If you make no modifications to the predefined profile values, the advanced mode will use the same settings as if you had selected the simple mode. When you have finished configuring the profile options, click **Next**.

Select the **Active Directory profile**. Click **Next** to continue.

8. Test Results

Directory Import Wizard

Test Results

11,112 users have been found.

Login Name	Name	DN	UID	GID	SID
user1	user1	CN=user1,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user2	user2	CN=user2,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user3	user3	CN=user3,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user4	user4	CN=user4,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user5	user5	CN=user5,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user6	user6	CN=user6,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user7	user7	CN=user7,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...
user8	user8	CN=user8,OU=AutoUsers,DC=test,DC=active-...			S-1-5-...

67 groups have been found.

Name	members	DN	GID	SID
Administrateurs de...	CN=Administrateur,CN=Use...	CN=Administrateurs de l'ent...		S-1-5-21-...
Administrateurs du...	CN=Administrateur,CN=Use...	CN=Administrateurs du sché...		S-1-5-21-...
Admins du domaine	CN=usermount usermount,...	CN=Admins du domaine,CN...		S-1-5-21-...
CERTSVC_DCOM_...	CN=Utilisa. du domaine,CN...	CN=CERTSVC_DCOM_ACCE...		S-1-5-21-...
Contrôleurs de do...		CN=Contrôleurs de domaine...		S-1-5-21-...
Duplicateurs		CN=Duplicateurs,CN=Builti...		S-1-5-32-...
Éditeurs de certific...	CN=WINUS,OU=Domain Co...	CN=Éditeurs de certificats,C...		S-1-5-21-...
Générateurs d'appr...		CN=Générateurs d'approbati...		S-1-5-32-...

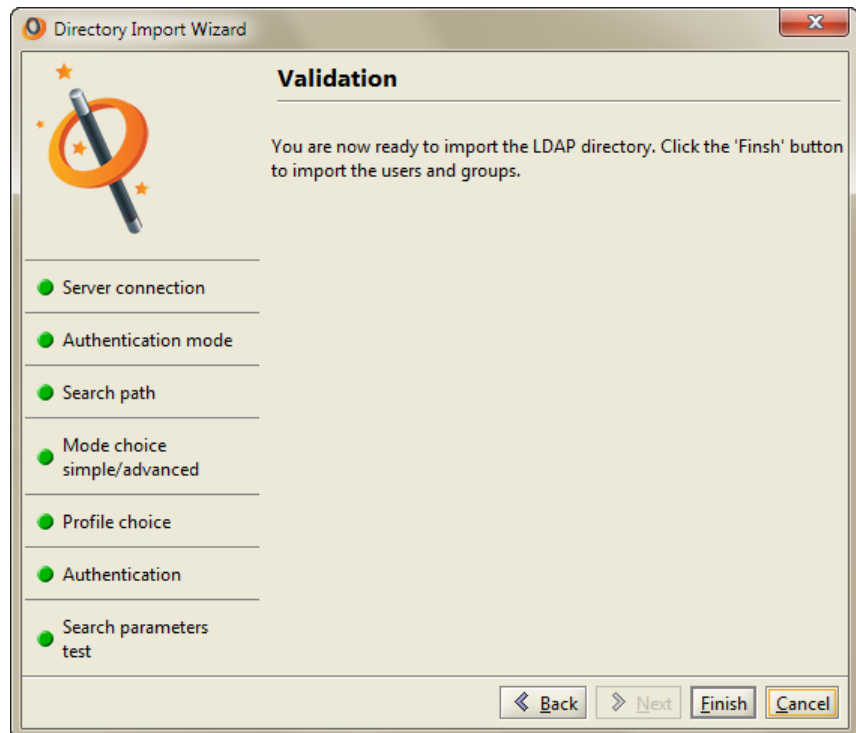
7 groups have been found for current user.

Name	members	DN	GID	SID
Administrateurs	CN=pad,CN=Users,D...	CN=Administrateurs,CN=Built...		S-1-5-32-544
Administrateurs de l'e...	CN=Administrateur,C...	CN=Administrateurs de l'entre...		S-1-5-21-312...
Administrateurs du sc...	CN=Administrateur,C...	CN=Administrateurs du sché...		S-1-5-21-312...
Admins du domaine	CN=usermount user...	CN=Admins du domaine,CN=...		S-1-5-21-312...
Propriétaires créateu...	CN=Administrateur,C...	CN=Propriétaires créateurs de ...		S-1-5-21-312...
SQLServer2005MSSQL...	CN=Administrateur,C...	CN=SQLServer2005MSSQLUser...		S-1-5-21-312...
Utilisateurs	CN=nouveauGroupe,...	CN=Utilisateurs,CN=Builti...		S-1-5-32-545

☐ Server connection
☐ Authentication mode
☐ Authentication
☐ Search path
☐ Mode choice simple/advanced
☐ Profile choice
☒ Search parameters test

If you have configured the import correctly, you should see the information from the external directory in the fields of the three tables for users, groups, and groups for the current user. This screen is only for information, so you can review the list of users and groups found in the directory. Click **Next** to continue.

9. Validation



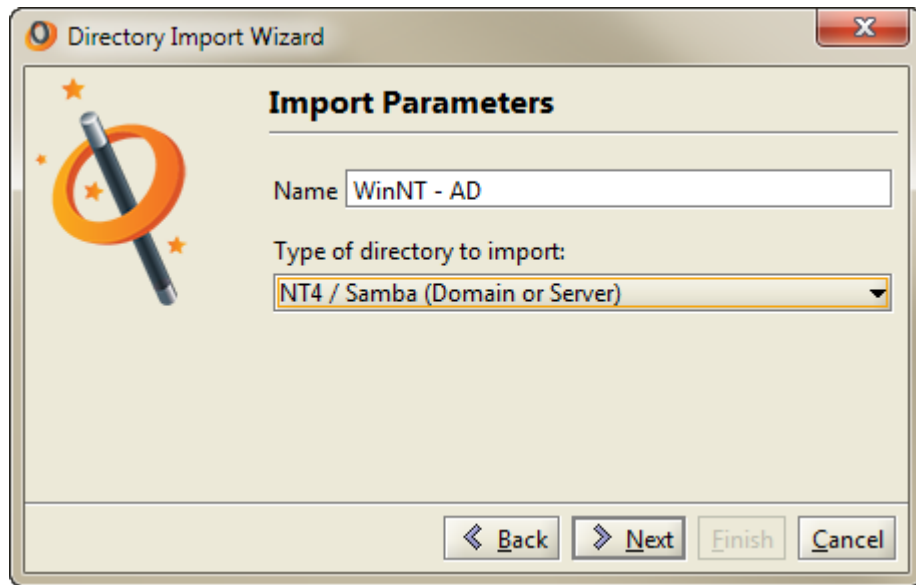
The configuration is complete. If you are ready to begin the import, click the **Finish** button. This begins the import process. The time required to perform this operation depends on the size of the directory being imported.

You can view or modify the connection information for the imported directory by selecting it in the **External Directories** tab and clicking the **Configuration** button.

Active Directory Imported in Windows NT Mode

This section describes the procedure to import an Active Directory in Windows NT Mode.

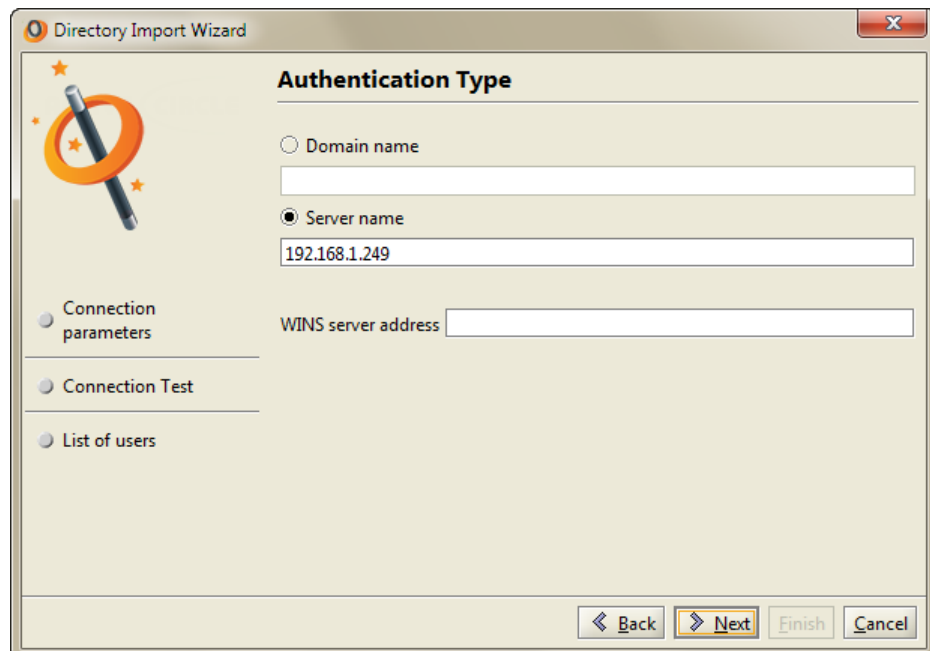
1. Import Parameters



The screenshot shows the 'Directory Import Wizard' window with the 'Import Parameters' tab selected. On the left is a graphic of a pen writing on a circle with stars. The main area contains a 'Name' text box with 'WinNT - AD' entered, and a 'Type of directory to import:' dropdown menu with 'NT4 / Samba (Domain or Server)' selected. At the bottom are four buttons: '< Back', '> Next', 'Finish', and 'Cancel'.

Type a **Name** that will be used to identify the imported directory withing Active Circle. Then select the option **NT 4 / Samba (Domain or Server)** from the drop-down list labeled **Type of directory to import**. Click **Next** to continue.

2. Authentication Type



The screenshot shows the 'Directory Import Wizard' window with the 'Authentication Type' tab selected. On the left is the same pen graphic. The main area has three radio buttons: 'Domain name' (unselected), 'Server name' (selected), and 'List of users' (unselected). Below 'Domain name' is an empty text box. Below 'Server name' is a text box containing '192.168.1.249'. Below 'List of users' is a text box labeled 'WINS server address'. On the left side of the main area are three sub-tabs: 'Connection parameters' (selected), 'Connection Test', and 'List of users'. At the bottom are four buttons: '< Back', '> Next', 'Finish', and 'Cancel'.

Specify the connection information for the server hosting the Active Directory. Depending on the configuration of your network, you must provide either:

- **Domain Name**
- **Server Name**

You may additionally enter the **WINS Server Address**, if applicable. Click **Next** to continue.

3. Authentication

Directory Import Wizard

Authentication

Login name: administrateur

Password: ••••••

Connection parameters

Connection Test

List of users

Please enter the login name and the password of a user with access rights on server.

Back Next Finish Cancel

Type the **Login name** and **Password** of a user who has access to the Active Directory. Click **Next** to continue.

4. Test Results

Directory Import Wizard

Test Results

38 users have been found.

Login Name	Name
user0001	user0001
user0002	user0002
user0003	user0003
user0004	user0004
user0005	user0005
user0006	user0006
user0007	user0007
user0008	user0008
user0009	user0009
user0010	user0010

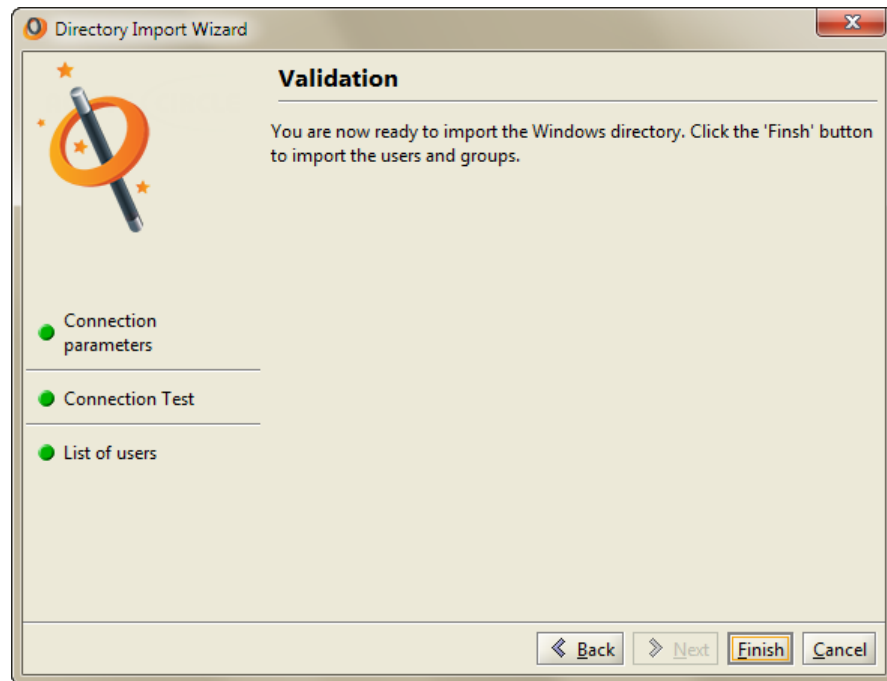
44 groups have been found.

Name	Members
groupOuGrp1	[user0010]
groupOuGrp1.1	[user0010]
groupOuGrp2	[user0010]
groupOuGrp2.1	
groupOuGrp3	
groupOuGrp3.1	
Invités du domaine	[Invité]
MonTest1	
MonTest2	

Back Next Finish Cancel

You should see the information from the external directory in the fields of the user and group tables. Click **Next** to continue.

5. Validation



You are now ready to import the Windows directory. Click the **Finish** button to begin the import process.

You can view or modify the connection information for the imported directory by selecting it in the **External Directories** tab and clicking the **Configuration** button.

NIS Directory

This section describes the procedure for importing users from a NIS directory.

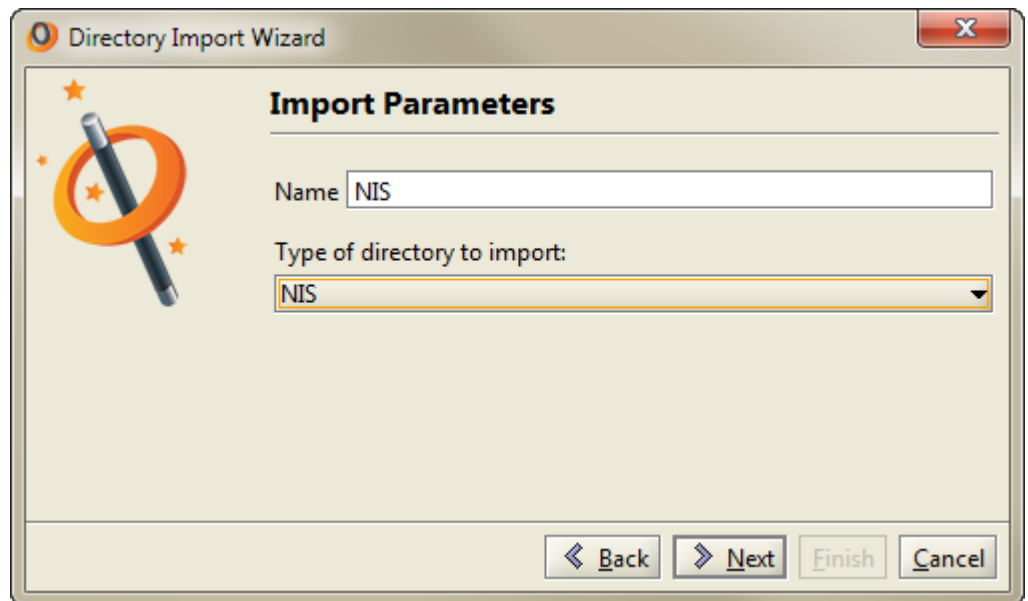


Important

The Active Circle name for a user imported from a NIS directory corresponds to the name provided in the `GECOS` field of the external directory.

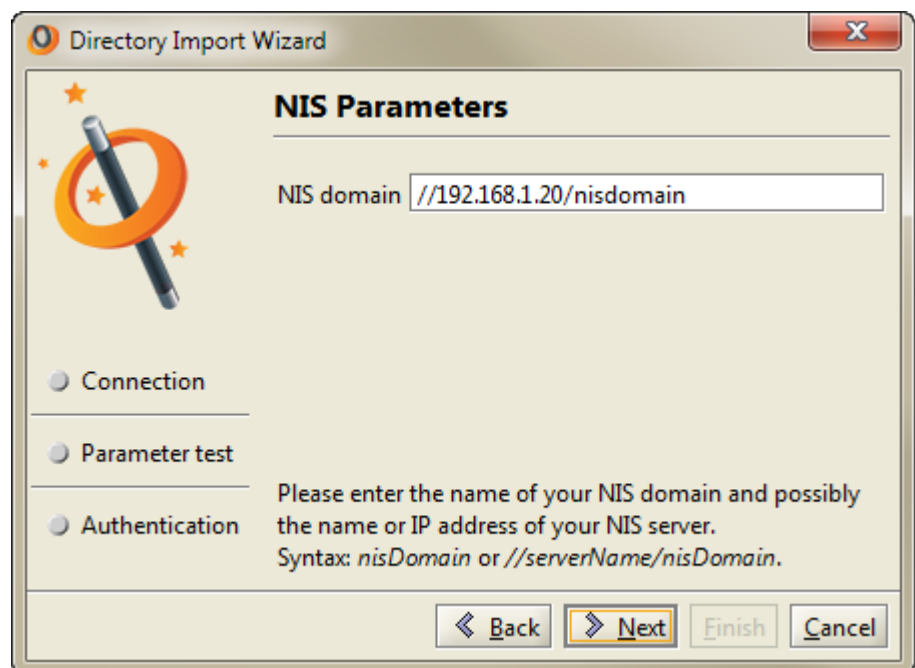
1. Import Parameters

Type a **Name** to identify the source of imported directory.



Select **NIS** from the drop-down list labelled **Type of directory to import**, and click on **Next**.

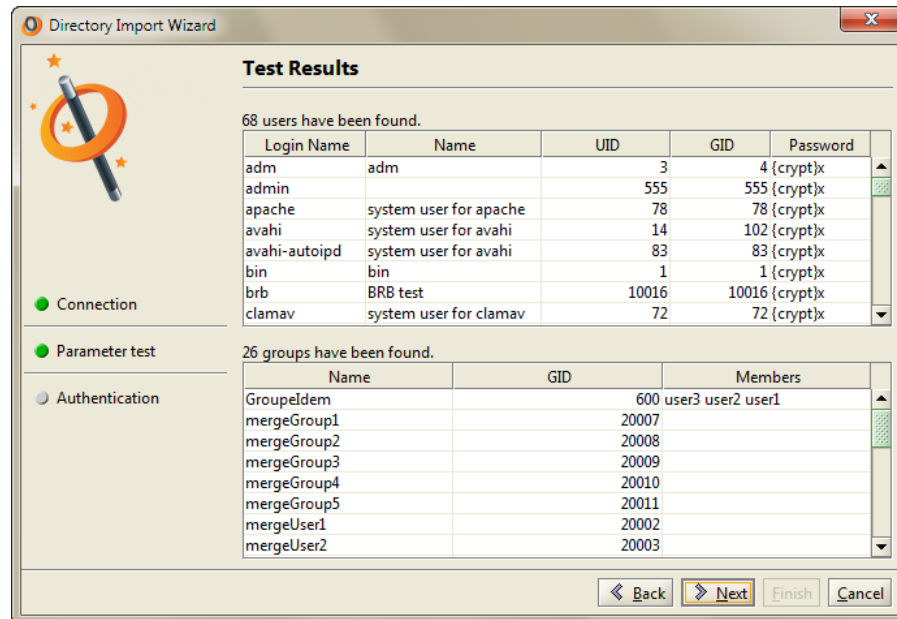
2. **NIS Parameters**



In this screen you must specify the path to the **NIS Domain**. When you click the **Next** button, Active Circle will try to connect to the directory.

3. **Test Results**

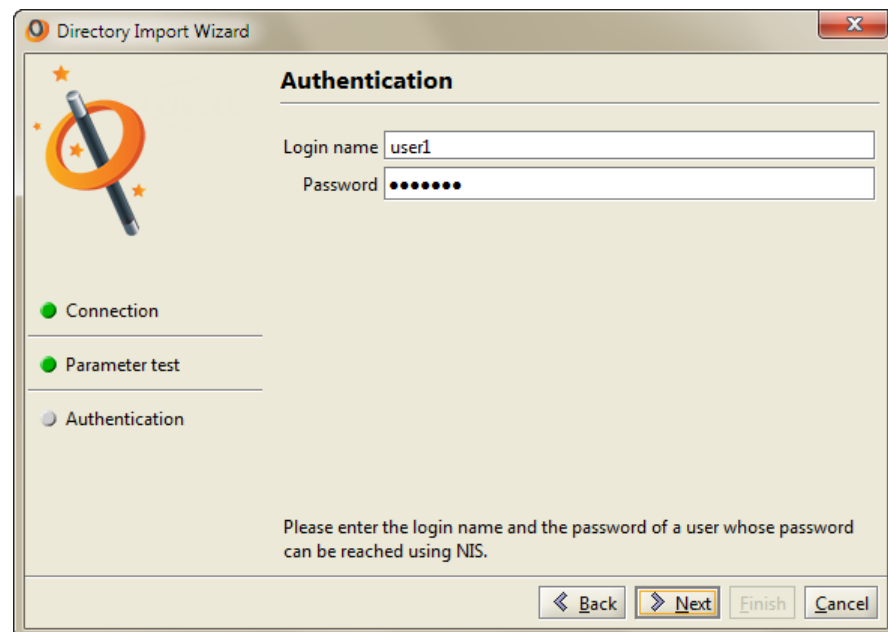
After Active Circle has queried the directory, the results are displayed in this screen.



The tables show the list of users and groups found in the directory. Click the **Next** button to continue.

4. Authentication


Enter the user name and password of a user with administrative rights on the domain.



Click the **Next** button to continue.

5. Validation

Click **Finish** to begin the import process.



You can view or modify the connection information for the imported directory by selecting it in the **External Directories** tab and clicking the **Configuration** button.

Chapter 21. Users and Groups

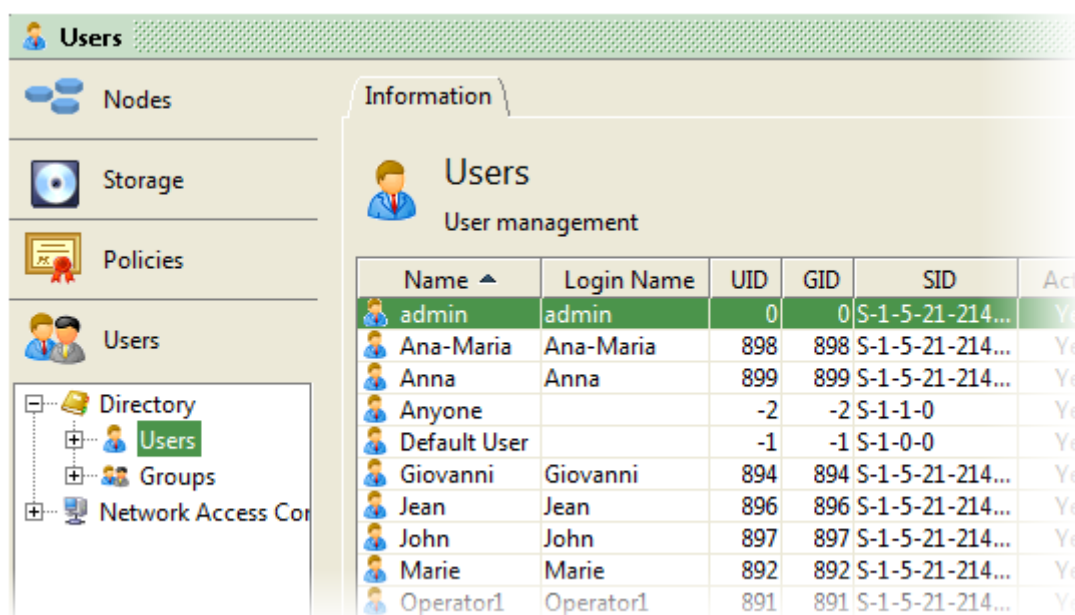
Use the **Users** and **Groups** items under **Directory** to manage user accounts and group memberships.

The tasks and settings for managing users and groups are very similar. In this manual, the process for managing **Users** also applies to **Groups** unless otherwise noted.

Viewing Users and Groups

Clicking on either the **Users** or the **Groups** items will display the list of users or groups in the **Information** tab to the right.

Figure 21.1. User List



The screenshot shows the 'Users' management interface. On the left is a sidebar with icons for Nodes, Storage, Policies, Users, and Directory. The 'Users' icon is selected. The main area is titled 'Users' and 'User management'. It contains a table with the following columns: Name, Login Name, UID, GID, SID, and Active status. The table lists several users, including 'admin', 'Ana-Maria', 'Anna', 'Anyone', 'Default User', 'Giovanni', 'Jean', 'John', 'Marie', and 'Operator1'.

Name	Login Name	UID	GID	SID	Active status
admin	admin	0	0	S-1-5-21-214...	Yes
Ana-Maria	Ana-Maria	898	898	S-1-5-21-214...	Yes
Anna	Anna	899	899	S-1-5-21-214...	Yes
Anyone		-2	-2	S-1-1-0	Yes
Default User		-1	-1	S-1-0-0	Yes
Giovanni	Giovanni	894	894	S-1-5-21-214...	Yes
Jean	Jean	896	896	S-1-5-21-214...	Yes
John	John	897	897	S-1-5-21-214...	Yes
Marie	Marie	892	892	S-1-5-21-214...	Yes
Operator1	Operator1	891	891	S-1-5-21-214...	Yes

By default, the table displays the following data for each user or group:

<i>Name</i>	<i>SID</i>
<i>Login name</i>	<i>Active status</i>
<i>UID</i>	<i>Directory source</i>
<i>GID</i>	

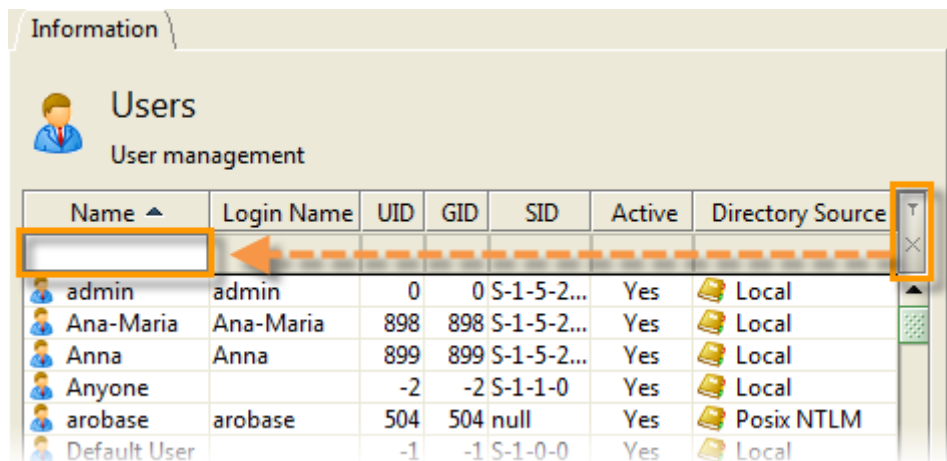
You can sort the list by any column by clicking on the column header. Click again to invert the sort order.

Searching

The user and group lists sometimes contain a large number of entries, making it difficult to quickly find the information you are looking for. It is therefore possible to search for specific users or groups by name.

To do this, click on the [filter button](#) to display the search field for the **Name** column.

Figure 21.2. User Filter

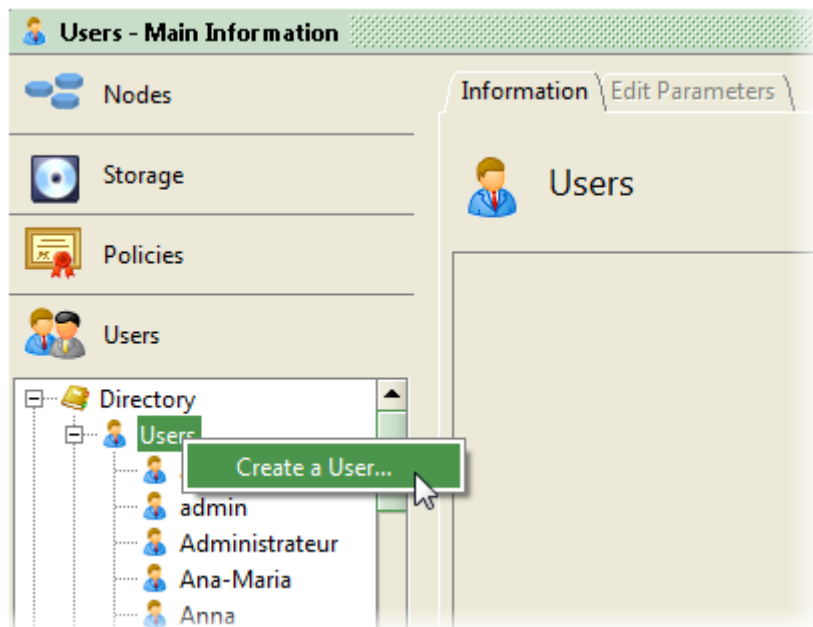


Type in any character or number in the search field. The list will be automatically updated as you type, to display only the entries that contain the characters you enter. For example, if you type the letter 'a', only the user names containing this letter will be listed. The filter is not case sensitive.

Create a Local User

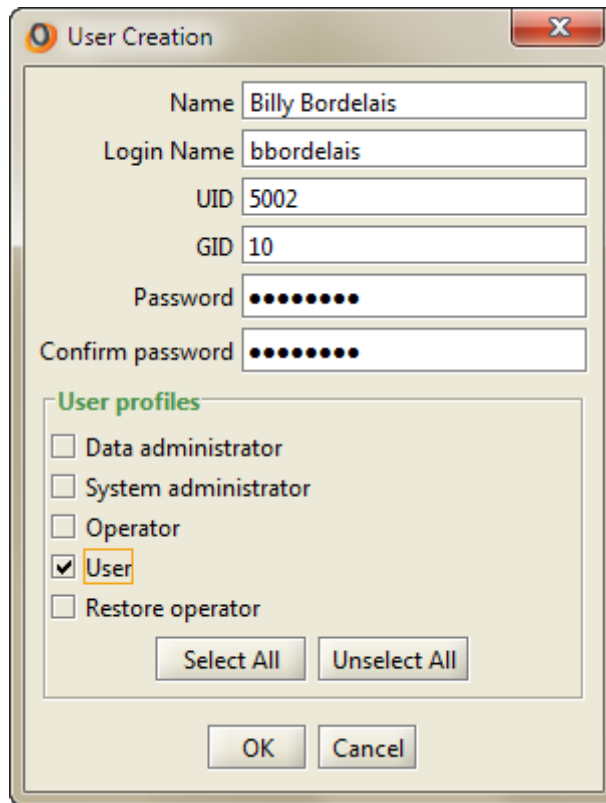
To create a new user, right-click on the **Users** item and select **Create a User....** To create a new group, right-click on the **Groups** item and select **Create a Group....**

Figure 21.3. User Creation



The **User Creation** dialog box appears with the fields required to create a new user:

Figure 21.4. User Creation — Enter User Details



The dialog box is titled "User Creation" and contains the following fields and options:

- Name:** Billy Bordelais
- Login Name:** bbordelais
- UID:** 5002
- GID:** 10
- Password:** (masked with dots)
- Confirm password:** (masked with dots)
- User profiles:**
 - ☐ Data administrator
 - ☐ System administrator
 - ☐ Operator
 - ☒ User
 - ☐ Restore operator
- Select All** and **Unselect All** buttons.
- OK** and **Cancel** buttons.

Name	A name to identify the user. <i>This is not the login name.</i>
Login Name	The name used to log in when accessing the system via CIFS or FTP.
UID	User Identifier. Used when accessing the system via NFS. It must be a unique integer value equal to or greater than 0.
GID	Group Identifier. Used when accessing the system via NFS. It must be an integer value equal to or greater than 0.
Password	Password used when accessing the system via CIFS or FTP. The password cannot be empty.



Note

The **Group Creation** dialog box does not have fields for **Login Name** and **UID**.

Next you must choose which **User profiles** to associate with the new user.

- **Data administrator**
- **System administrator**
- **Operator**
- **User**
- **Restore operator**

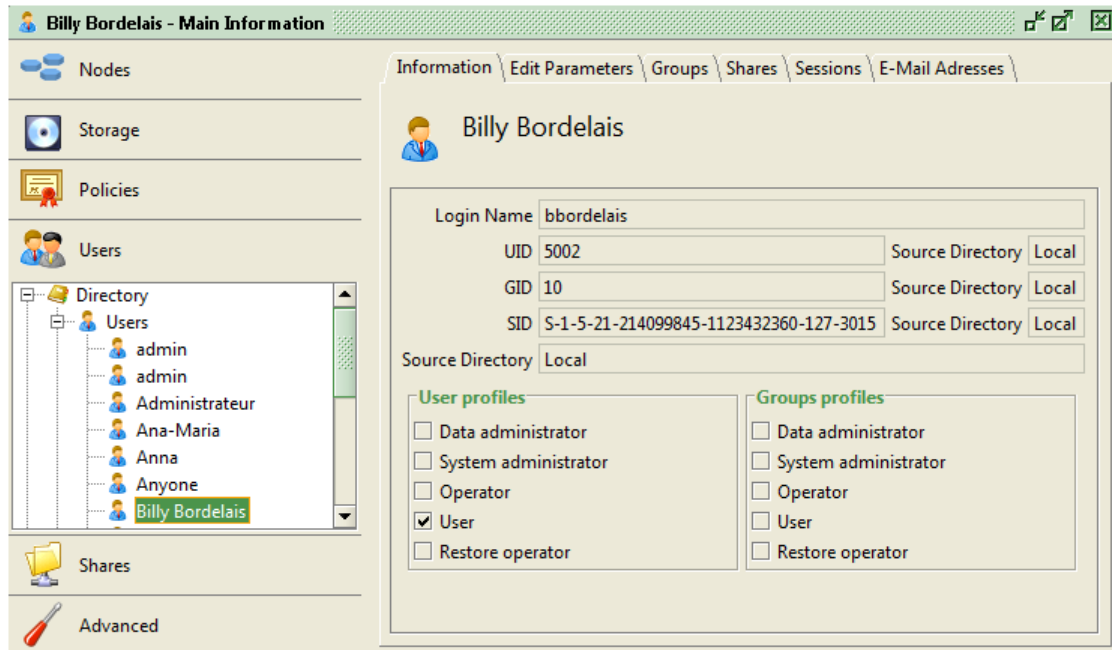
For more information about **User profiles**, please see ["Profiles"](#).

The buttons **Select All** and **Unselect All** allow you to select or to unselect all the **User profiles** respectively.

Click **OK** to create the user or **Cancel** to abort user creation.

After the user has been created, you can view the user details by selecting the item in the **Users** view:

Figure 21.5. User Created



By right-clicking on a user name, a menu is displayed where you can **Deactivate** or **Delete** the user. Groups can be deleted in the same way.

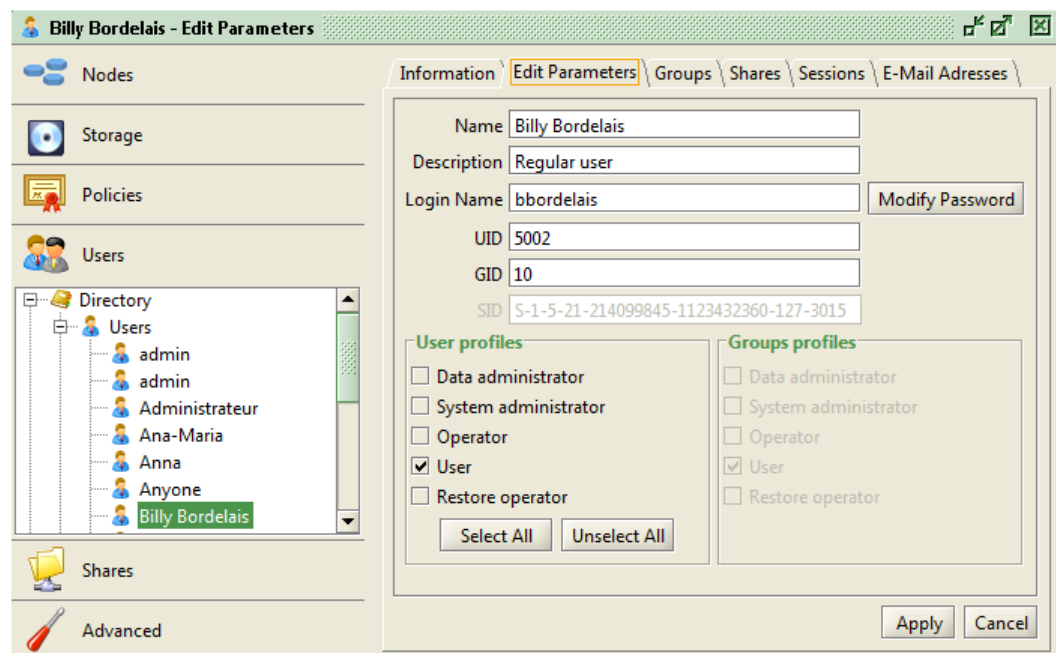
Information

The **Information** tab provides an overview of the user account. The field **Source Directory** indicates the directory service the user was created in. The value **Local** indicates that it is an Active Circle user. If the user was created in an external directory, OpenLDAP for example, the name of the external directory would be displayed.

Modifying User Settings

The **Edit Parameters** tab allows you to modify information relating to the selected user or group.

Figure 21.6. Edit Parameters tab



The fields and settings you can modify on this tab are mostly the same as during [user creation](#). In addition, you can provide a description for the user.

The **SID** (Security Identifier) cannot be modified. This is used by Microsoft Windows ACLs to identify users and their group memberships in order to grant or deny access to resources.

Changing User Profile

There are two sections that display the profiles associated with the user account:

- User profiles** The profiles assigned directly at the user level. A profile can be added or removed here by checking or unchecking it. You must click the [Apply](#) button to validate any changes.
- Group profiles** Greyed out. These options display the profiles inherited by the user from any **Groups** the user belongs to.

You can use the [Select All](#) and [Unselect All](#) buttons to facilitate associating **Profiles** with user accounts.

For more information about the **Profiles**, please see [“Profiles”](#).

Changing User Passwords

Use the [Modify Password](#) button to change the password associated with the Active Circle user account.



Important

The password for the user `admin` should be changed immediately after installation.

Managing Groups

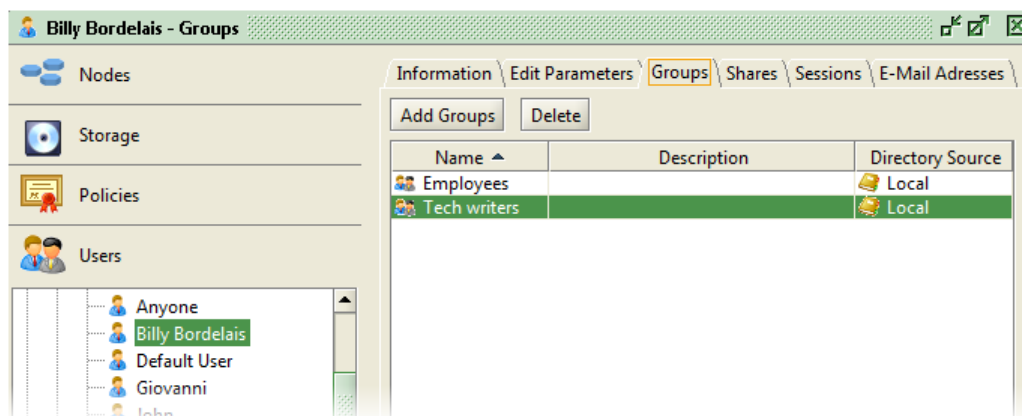
You can add and remove users to and from groups either from the perspective of a user or that of a group. If you select a *user name* from the list in the **Users** view, you can add or remove group memberships for the user by using the **Groups** tab. If you select a *group* from the list in the **Users** view, you can use the **Members** tab to add or remove users from the group. The group item also has a **Parent Groups** tab, which you can use to include a group in another group. All three tabs have the same layout; a list of users or groups, and an **Add** and a **Delete** button.

For all the group management tasks described, you can select users and groups from both the local and the external directories.

User Tasks

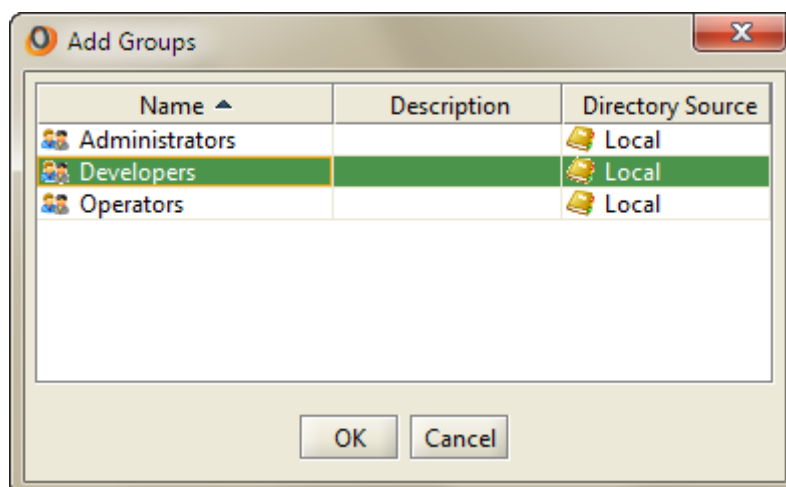
The **Groups** tab — which only appears under the **Users** item — allows you to manage the **Groups** to which the user belongs.

Figure 21.7. Groups tab



To add a user to a **Group**, click the **Add Groups** button. The **Add Groups** dialog box appears.

Figure 21.8. Add Groups dialog box



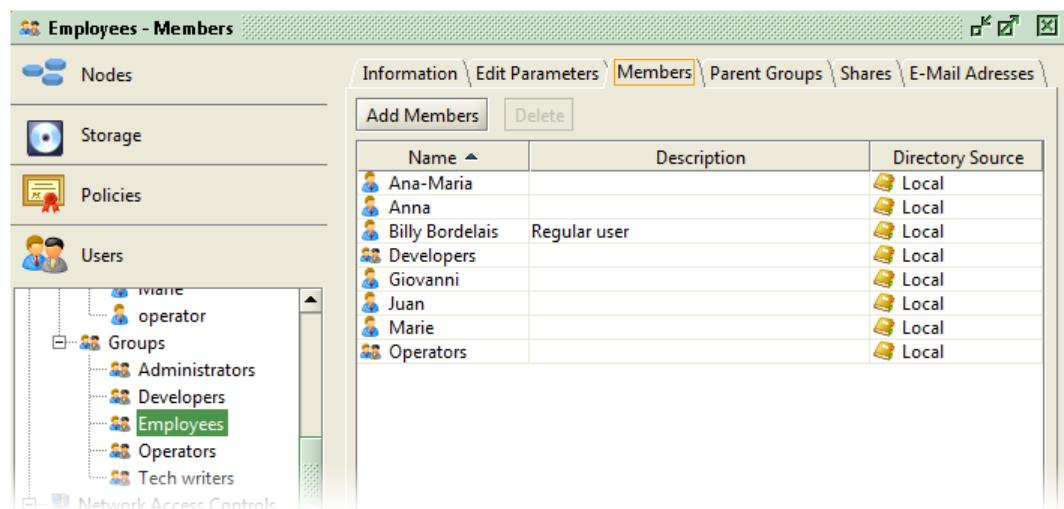
Select a group from the list and click **OK** to add the **Group**. You can add both local groups and groups from external directories.

You can remove a user from a **Group** using the **Delete** button. Note that this only removes the user from the group, it does not delete the group.

Group Tasks

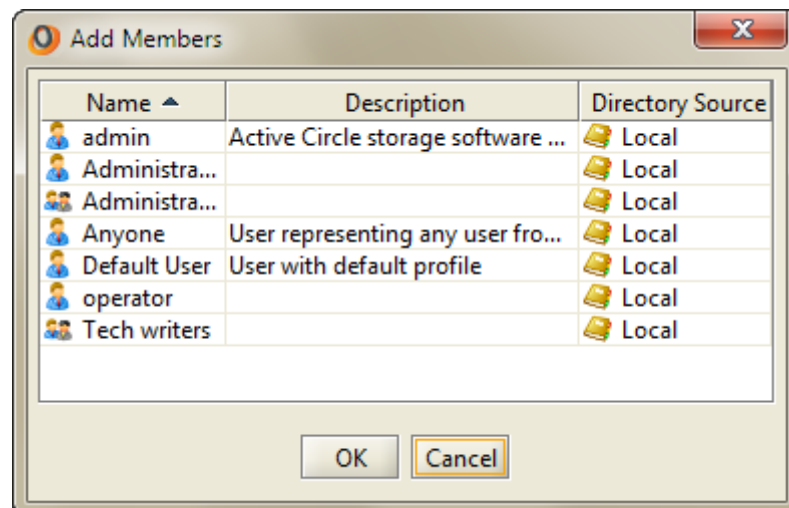
The **Members** tab — which only appears under the **Groups** item — allows you to manage user group memberships from a group perspective.

Figure 21.9. Members tab



To select one or more users or groups to add to the group, click the **Add Members** button.

Figure 21.10. Add Members dialog box

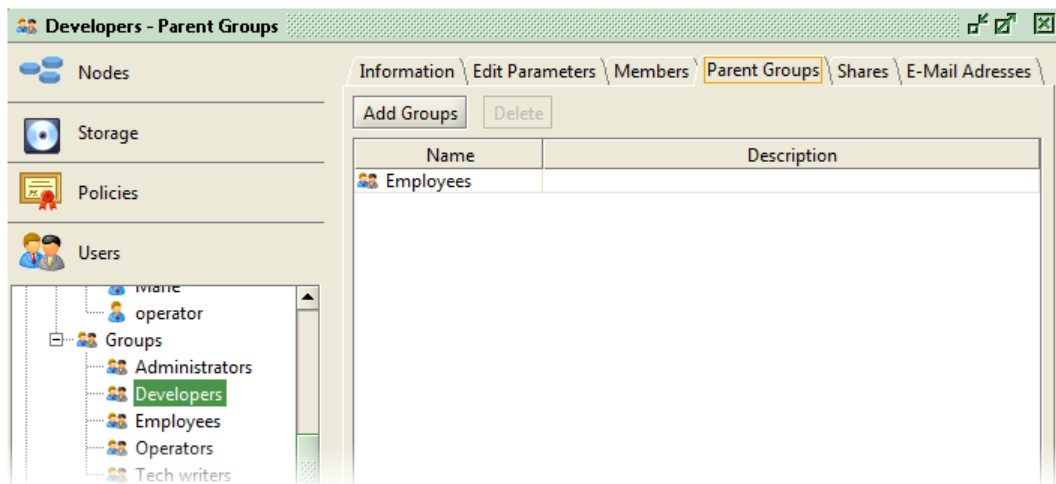


You can remove a user from the group using the **Delete** button.

Parent Groups

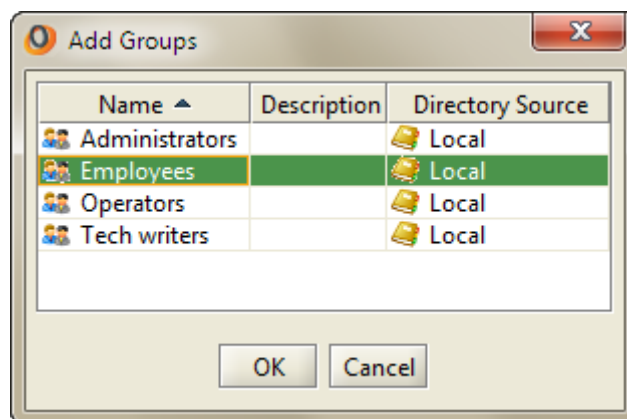
The **Parent Groups** tab — which only appears under the **Groups** item — allows you to create hierarchies of groups so that the profiles and access to Active Circle shares are inherited by child groups from their parents.

Figure 21.11. Parent Groups tab



To add a parent to the current group, click the **Add Groups** button. A dialog box appears in which you can select the new parent. Click **OK** to close the dialog box and add the user.

Figure 21.12. Add Members dialog box

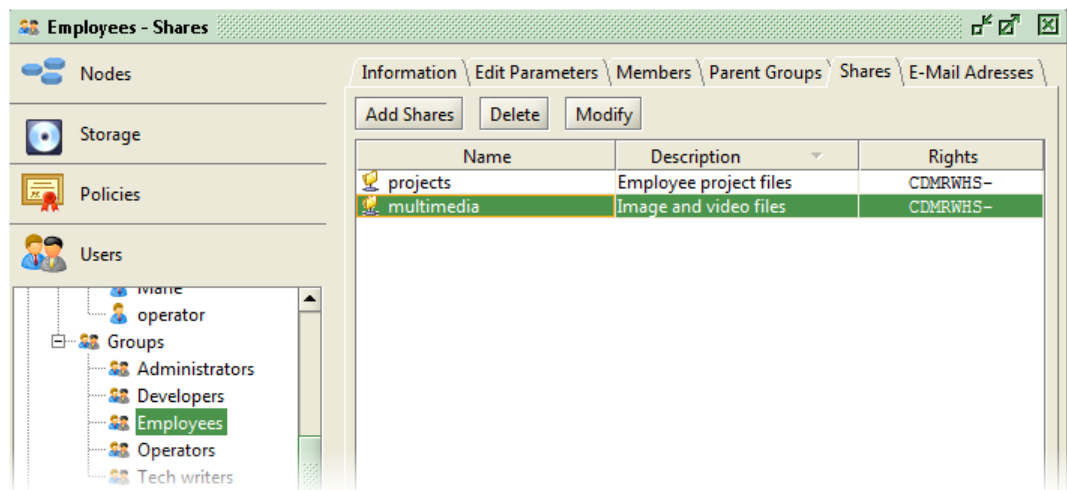


You can remove a parent group using the **Delete** button.

Share Access

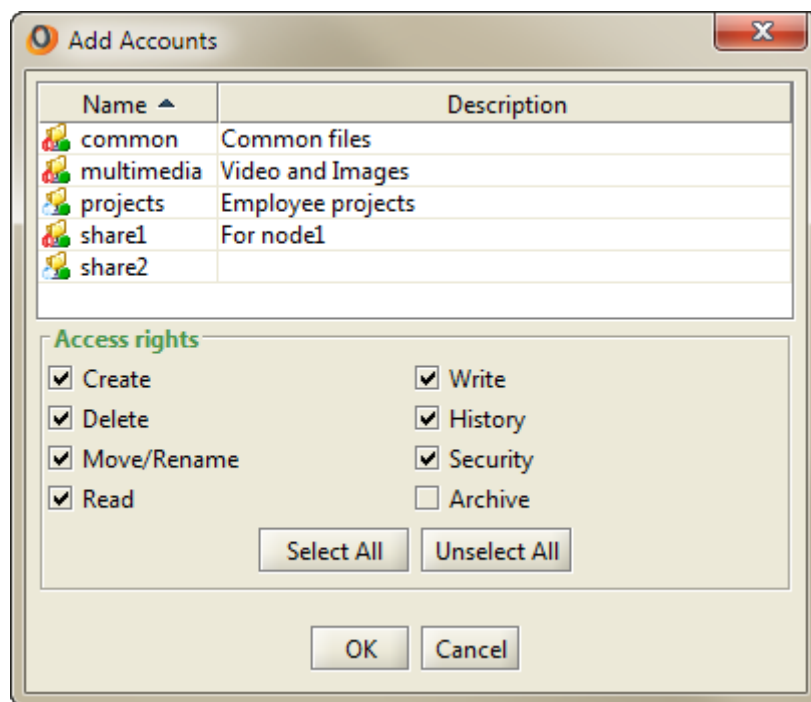
In the **Shares** tab, which exists on both the *user* and *group* levels, you can view and manage user and group access to shares.

Figure 21.13. Shares tab for a Group



To add a share, click the **Add Shares** button. The **Add Accounts** dialog box appears.

Figure 21.14. Add Accounts dialog box



At the top of the dialog box is a table listing the shares available for selection. Below that is an area displaying the **Access Rights** to be granted to the user or group on the share.

Table 21.1. Share Access Rights

Permission	Actions allowed
Create	Create new files and folders on the share
Delete	Delete files and folders from the share
Move/Rename	Move and/or rename files and folders on the share

Permission	Actions allowed
Read	Read files and folders on the share
Write	Write to files and folders on the share
History	Recover previous versions of files on the share
Security	Modify security settings for files and folders on the share
Archive	Initiate an on-demand archiving process for the share using the Active Circle File Explorer



Note

The **Archive** access right, which grants users the right to generate on-demand archives, is unchecked by default. For more information about on-demand archiving, please see the *Active Circle File Explorer Guide*.

To add user or group access to a share, select the share and choose the access rights to apply, and click **OK**.

The new share appears in the list. The **Rights** column displays the user's individual access rights on a given share, whereas the **Cumulative Rights** column displays the sum of user permissions coming from both individual and group assignments. If a user receives all the access rights for a share through a group membership, the **Rights** column will be empty. The **Shares** tab for a group does not have a **Cumulative Rights** column.

Figure 21.15. Shares tab for a User

Name	Description	Rights	Groups	Cumulative Rights
multimedia	Image and vide...		Employees	CDMRWHS-
projects	Employee proje...		Employees	CDMRWHS-
share-01	Main data share	CDMRWHS-		CDMRWHS-

You can modify the access rights for the user or group on the share using the **Modify** button.

Use the **Delete** button to remove the user's or group's access to the share.



Note

Deleting a share from a user does not necessarily remove all of a user's access rights to that share. If the user has been granted access in the context of a group, the user will still have access to the share. Individual rights will be removed (reflected in the **Rights** column), but the share will remain visible in the table as long as the user has access to the share in the context of the group.

User Sessions

The **Sessions** tab provides an overview of a user's activity in the system.

Figure 21.16. Sessions tab

Login Date	Type	Client Address	Last Activity	Node	Active	Begin	End
August 23, 2011 12:11:06 PM	Exploration	192.168.1.50	August 23, ...	node-01	No		
August 23, 2011 10:01:24 AM	Administration	192.168.1.50	August 23, ...	node-01	Yes		
August 23, 2011 9:31:03 AM	Administration	192.168.1.50	August 23, ...	node-01	No		
August 23, 2011 9:19:24 AM	Operation	192.168.1.50	August 23, ...	node-01	No	Janu...	Janu...
August 23, 2011 9:17:20 AM	Operation	192.168.1.50	August 23, ...	node-02	No	Janu...	Janu...
August 23, 2011 9:15:07 AM	Administration	192.168.1.50	August 23, ...	node-01	No		
August 22, 2011 6:17:16 PM	Operation	192.168.1.50	August 22, ...	node-01	No	Janu...	Janu...
August 22, 2011 2:32:40 PM	Operation	192.168.1.50	August 22, ...	node-01	No	Janu...	Janu...
August 22, 2011 2:32:39 PM	Operation	192.168.1.50	August 22, ...	node-02	No	Janu...	Janu...
August 22, 2011 12:44:14 PM	Administration	192.168.1.50	August 22, ...	node-01	No		
August 22, 2011 12:42:38 PM	Administration	192.168.1.50	August 22, ...	node-01	No		
August 22, 2011 12:13:08 PM	Administration	192.168.1.50	August 22, ...	node-01	No		
August 22, 2011 12:09:45 PM	Administration	192.168.1.50	August 22, ...	node-01	No		

The **Close** button allows an administrator to terminate a user session.

The following data columns are displayed by default:

Login Date

Last Activity

Type

Node

Client Address

Active

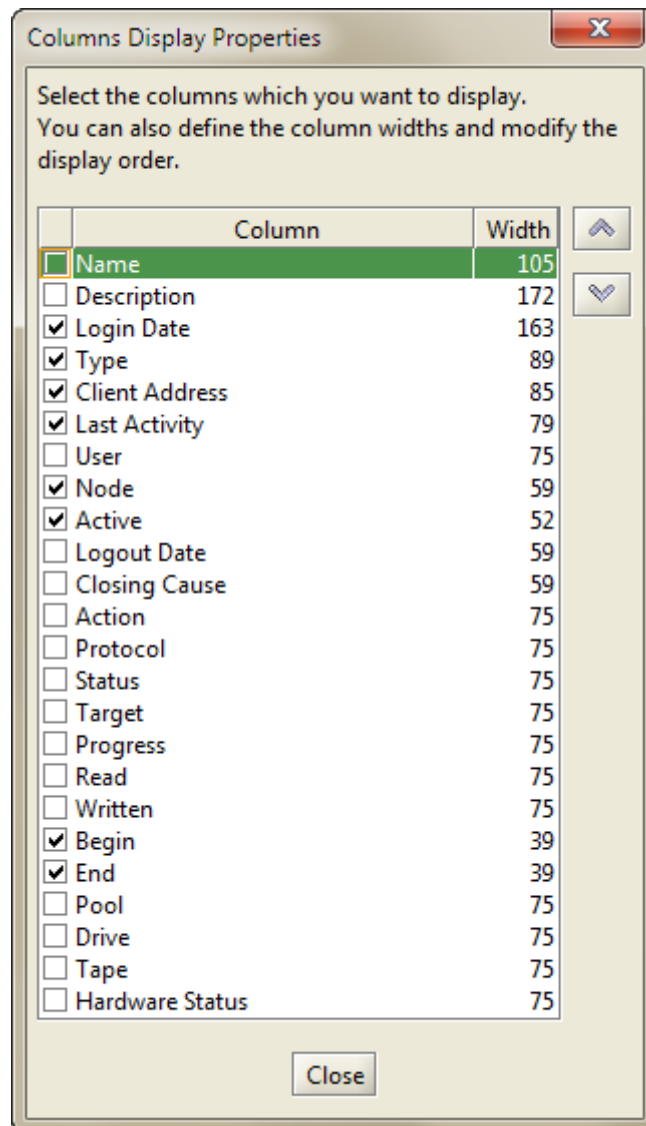
To display more (or fewer) columns of data, and to specify the order and width of the columns, *right-click on a column heading* and select **Column properties...** from the context menu being shown. This will give you a list of many additional log data fields to display.



Note

To view details and progress information for *Operation* sessions, you must select more columns for display. For example, if you have activated **FTP session monitoring** by setting the *circle property* `session.monitored` to `true`, you should select columns like **Action**, **Protocol**, **Status**, **Progress**, **Read** and **Written** to be able to see the progress and details of the session.

Figure 21.17. Columns Display Properties Dialog



Viewing Session Logs

The **View Logs** button allows you to open a reporting tool to obtain more detailed information on the user's activity.



Note

The logs will only be available if this functionality has been enabled. Due to performance considerations, session log creation is deactivated by default. You can activate it by setting the *circle property* `session.log` to *true*. For more information on how to set values for circle properties, see ["Circle Properties"](#).

The following log details are displayed in the viewer by default:

Date

Log Service Name

Thread ID

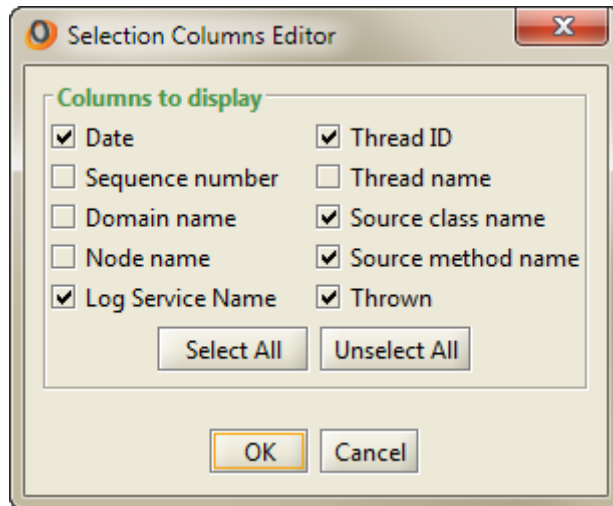
Source class name

Source method name

Thrown

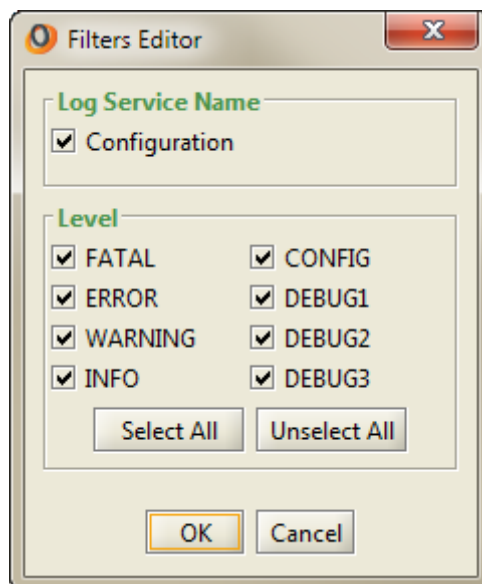
You can add or remove data columns from the viewer by clicking on the **Select Columns** button:

Figure 21.18. Selection Columns Editor



You can filter the log display by clicking on the **Filters** button. A dialog box will be displayed that lets you filter by **Log Service Name** and log entry importance level (information, warning, error, debug, etc.):

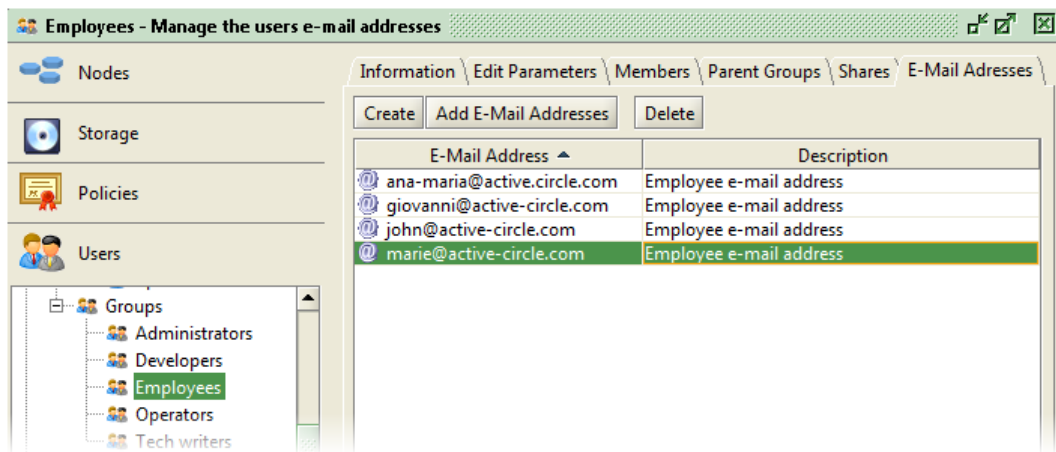
Figure 21.19. Filters Editor



E-Mail Addresses


This tab allows you to associate one or more e-mail addresses with a user account or a group. Active Circle can then generate e-mail notifications on certain events and send them to the appropriate user(s).

Figure 21.20. E-Mail Addresses tab for a Group



To add an e-mail account to a user, click the **Create** button.

Figure 21.21. Create an E-Mail Address tab



You can associate existing e-mail addresses with the user or group by clicking the **Add E-Mail Addresses** button. You can then select one or more addresses from a list.

Use the **Delete** button to remove an e-mail address from a user.

Chapter 22. Network Access Controls

In addition to assigning access rights and profiles to users and groups, you can also define access rules for your network.

For example, you could define a rule for a subnet that limits access to the system to user-level actions. By adding the Network Access Control to a share and granting read-only access, you can ensure that any users coming from that particular part of the network will only be able to consult (not modify) the resources stored in the share.

There are two types of Network Access Control:

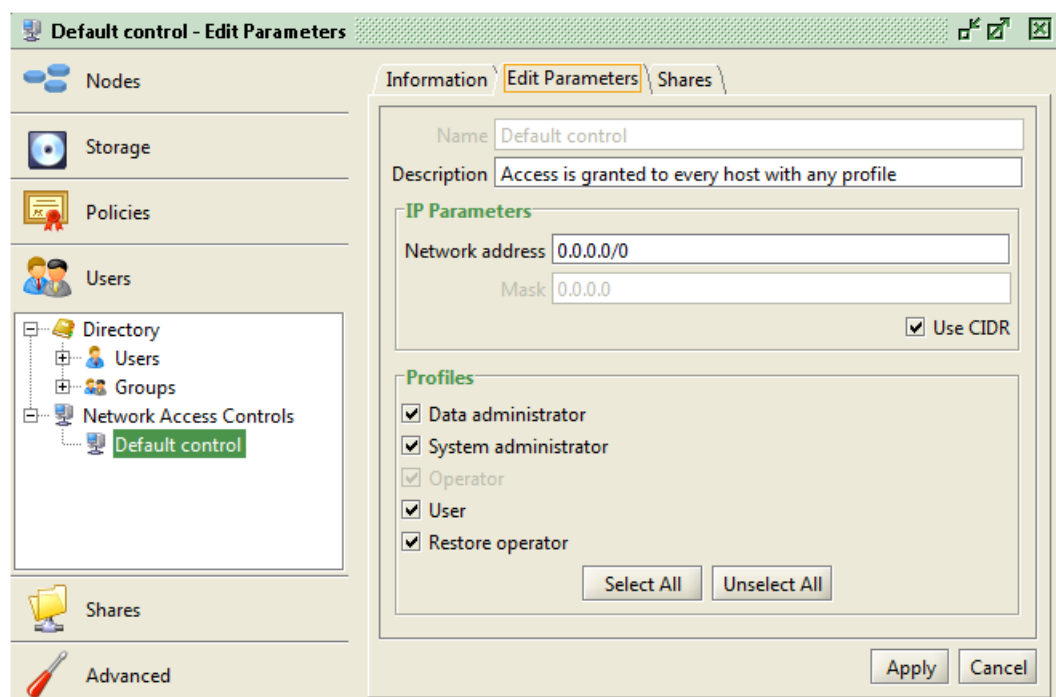
- [IP Network Access Control](#)
- [Host Name Network Access Control](#)

Each of these will be described in the sections below.

Default Control

When you install the first node, a **Default Control** is automatically created and configured. It is set to grant access to every host on the network with any profile. For this control to have any effect, it must be associated to an Active Circle share. When applied to a share this Network Access Control allows any user from any IP address to access the system. Whether you use this control or whether you apply a more specific level of control to your system depends on the context in which the Active Circle software is used.

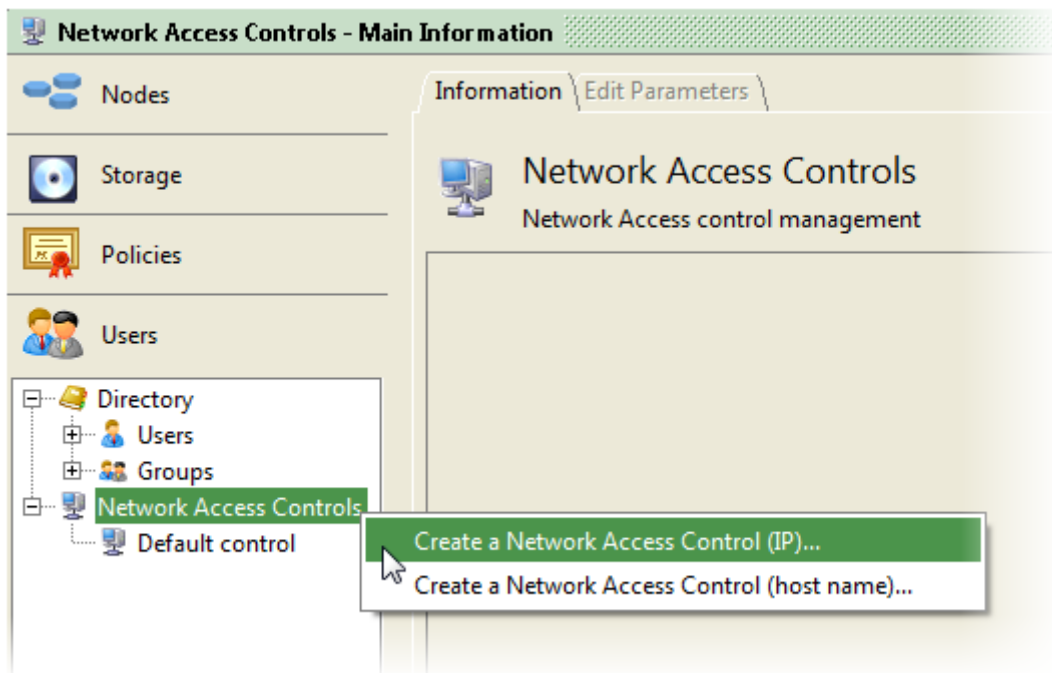
Figure 22.1. Default Control Settings



Create an IP Network Access Control


To create an IP Network Access Control, right-click on the Network Access Control item and select **Create a Network Access Control (IP)**.

Figure 22.2. Create a Network Access Control (ip)



The **Network Access Control** creation dialog box appears.

Figure 22.3. Network Access Control Creation



You must first provide a **Name** to identify the new Network Access Control. Next you must provide the **IP Parameters** for the control.

Defining IP Parameters

You can specify the network address in CIDR notation using the IP address and a routing prefix, or as an IP address with a subnet mask. The input fields will automatically convert the **Network address** and **Mask** fields when **Use CIDR** is selected or deselected (the conversion takes a few seconds, and the input fields and check box will not be available during this time).

Network address	The IP address to which the control will apply. In the example above, all access from the IP address 192.168.121.151 is limited to user actions only (the routing prefix /32 indicates all the bits of the address should be applied). If you deselect Use CIDR , the address will be converted to regular dot-decimal IP address notation (the /32 routing prefix will be removed), and the corresponding subnet mask (255.255.255.255) will appear in the Mask field.
Mask	The network mask to be applied to the Network Address . This field is greyed-out if the Use CIDR check box (see below) is activated. To specify the mask for the IP address from the example above without using CIDR, you would fill in 255 . 255 . 255 . 255 (meaning the mask is limited to the specific machine's address).
Use CIDR	Use the <i>Classless Inter-Domain Routing</i> method. For example, if you want to include all the machines in the 192.168.1.0 subnet, you specify 192.168.1.0/24 in the Network address field and select Use CIDR . Turning this option off activates the Mask field and converts the input fields into the corresponding regular address (192.168.1.0) and subnet mask (255.255.255.0) notation. For more information on CIDR, see this Wikipedia article .

Selecting User Profiles

In the area labeled **Network Access Control profile** you must set which profiles are allowed to use this Network Access Control.

- **Data administrator**
- **System administrator**
- **Operator**
- **User**
- **Restore operator**

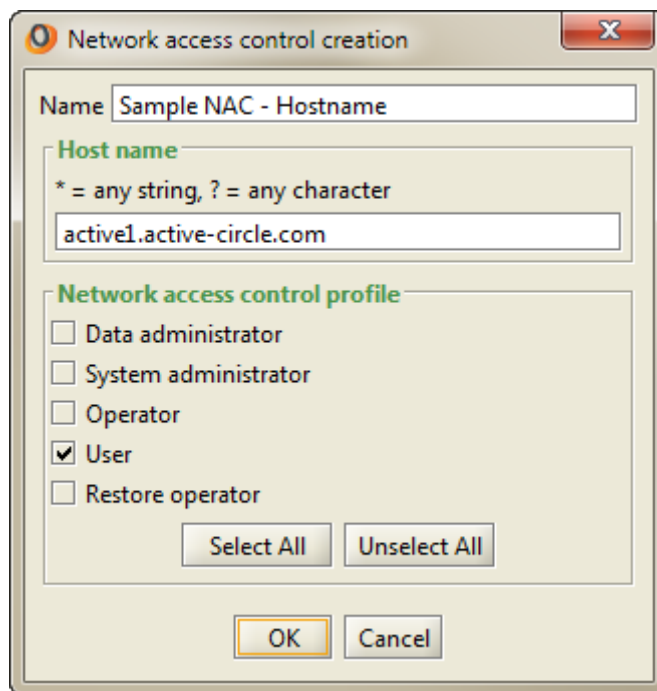
For more information about **User profiles**, please see [Chapter 18, User Profiles and Access](#).

Create a Host Name Network Access Control

In networks using DHCP, using Network Access Controls based on IP addresses might not be possible. The alternative is to use Network Access Controls based on the names of machines. With this option, it is possible to define rules using wildcard characters.

To create a host name Network Access Control, right-click on the Network Access Control item and select the option **Create a Network Access Control (host name)**.

Figure 22.4. Network Access Control creation (host name) dialog box



Provide a **Name** and a **Host Name** for the new control. For the host name, you may either enter the exact name of the machine for which you wish to define the rule, for example `server1.active-circle.com`; or you may use one or both of the wildcard characters "*" or "?" to represent multiple machines, for example `server?.active-circle.com` (* can represent any string of characters, while ? represents any single character). For all protocols, both the short name and the fully qualified domain name are valid for network access control by Active Circle.

**Note**

The **Host Name** field is not case sensitive.

Select the profile(s) to apply to the control. For more information about profiles, please see ["Profiles"](#).

Apply a Network Access Control to a Share

For the Network Access Control to have an effect on system access, you must associate it with one or more shares, using the **Shares** tab. There you click on the **Add Shares** button to add one or more shares and define the access rights. The procedure and options are the same as when associating a user or group with a share (see ["Share Access"](#) in [Chapter 21, Users and Groups](#)), using the **Add Account** dialog box.

Alternatively, you can add a network access control to a share using the **Access Controls** tab of a selected share in the **Shares view**. Click on the **Add Network Access Controls** button to add one or more previously defined network access controls to the selected share.

Chapter 23. Shares


A Share provides the users with a point of access to the files stored in the Active Circle Virtual File System (VFS). An Active Circle share functions in a similar way to a shared directory in Microsoft Windows. By mapping a drive to the share's path, a user can browse the directory structure of a share and access the files stored there in the same way as if the files were stored under a traditional network share. For most users, their only contact with the Active Circle Storage System would be through a mounted Active Circle share, which would look like any other mapped network drive or mounted data share in the operating system of the user.

In the Administration Tool, you can get an overview of all the shares in the system and perform basic tasks on the **Shares** tab of the root **Shares** item.

Figure 23.1. Share List

Name	Group	Numb...	Volume	Policy ...	Storage Polic
multimedia	User Data	40	403.04MB	node-02	1 x per domain
share-01		23,422	2.29GB	node-02	1 x per domain
test		46	613.56MB	node-01	1 copy - 1 copy
projects	User Data	1,155	1.21GB	node-01	1 x per domain

The **Information** box at the top shows summary information for all the shares of the system. The table below it lists the shares, including many columns of data relevant to each share, such as the [policy status](#), which includes the number of *obsolete files* on each share. You can use the buttons above the share list to add, delete or move shares. These operations are described later in this chapter.

**Note**

The number of shares in the system is limited to 100 by default. You can increase or decrease this limit by modifying the value of the [circle property](#) `vfs.maxShareCount`.

Getting Started

Shares is a central concept in the Active Circle architecture, as they are the users' point of contact with the system. Many aspects of the system configuration are associated with

managing shares, which means that some planning should be involved in creating and configuring shares. The list below summarizes the main steps that need to be performed to make shares work, indicating the associated components.

General steps for preparing and using an Active Circle share:

- Create a policy (see [Chapter 15, Storage Policy](#))
- [Create a Share](#) and associate it with a policy
- Add users and define [Access Controls](#) for the share
- [Activate the share on one or more nodes](#)
- Activate the file server on the nodes (see “[Activating the File Server](#)” in [Chapter 24, Protocols](#))
- [Connect to the share](#) using CIFS (Windows), NFS or FTP
- [Deposit or migrate files to the share](#)

In addition to these general tasks which must be completed, there are several optional share and file server configuration tasks which are also described in this (and the next) chapter.

Create a Share

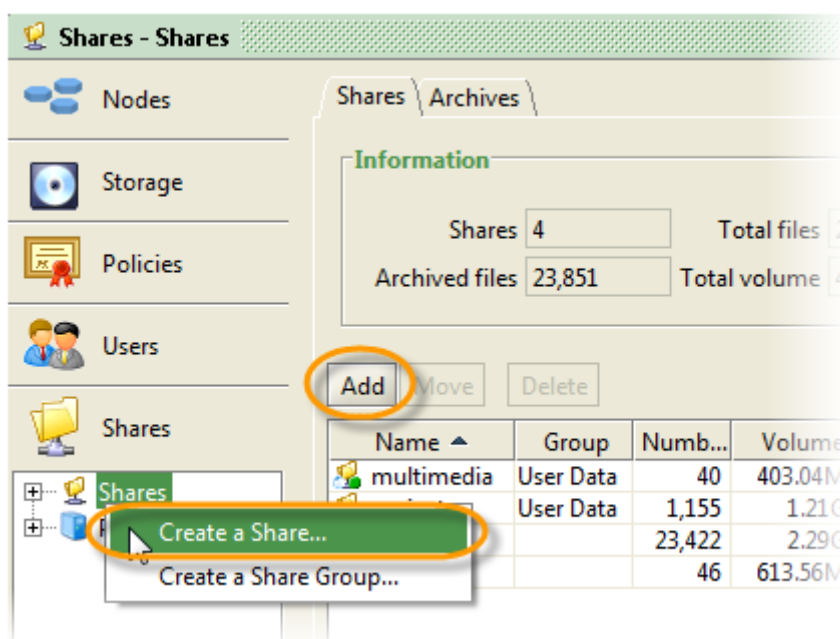


Important

All shares need to be associated with at least one storage or archiving policy. If one does not already exist, you will be asked to create one as part of the share creation process. For more information on creating policies, see “[Create a Storage Policy](#)” or “[Create an Archiving Policy](#)”.


To create a new share, right-click on the **Shares** item in the **Shares View** and select **Create a Share....** Alternatively, you can click on the **Add** button on the **Shares** tab.

Figure 23.2. Create a Share



Enter the name for the new share in the dialog box and select the storage policy and/or Archiving Policy to be associated with it. We recommend using only lowercase letters in the share name, and no spaces. The following special characters are not allowed: \ / * ? " < > |.

Figure 23.3. Create a Share: Selecting a Policy



If no suitable policy exists in the lists, you can create a storage and/or an archiving policy from this dialog by clicking on the respective **Create** buttons to the right above each policy list. This will open the regular policy creation dialogs, described in the **Policies View** part of this guide.



Note

A policy can be associated with more than one share.

Once you have entered a name and selected a policy, click **OK** to create the share. The new share appears in the list under the **Shares** element. It will now also appear under the **Shares** tab for the policy selected during creation in the list of associated shares, as well as on the **Shares** tab for each node.

Certain share properties are automatically configured by the system upon share creation:

- The new share is automatically activated on all the nodes of the circle.
- The default user `admin` is added to the **Access Controls** with all access rights.
- The default quota for the share size is set to 1 petabyte.
- Storage and versioning policy processing is activated by default. Archive policy processing is not automatically activated.
- Choice of policy processing node is set to **Automatic**.

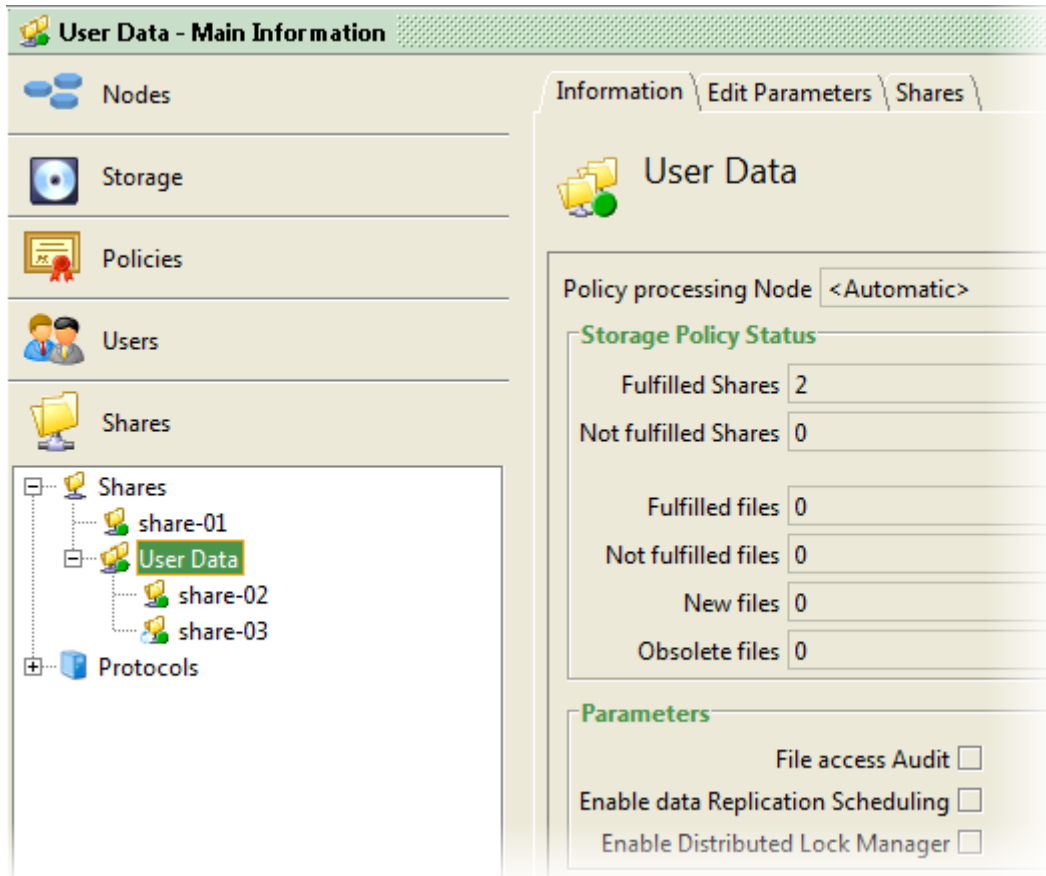
Create a Share Group

You can also create Share Groups in order to simplify the administration of the various shares managed by Active Circle.

To create a Share Group, right-click on the **Shares** item in the **Shares View** and select **Create a Share Group....** Enter a name for the share group and click **OK** to create the group.

You can create shares directly in the group or move them to a group using the **Move** function (see [below](#)). In the example given below, there are two shares — `share-02` and `share-03` — in the share group `User Data`.

Figure 23.4. A Share Group



Each Share Group has its own **Information**, **Edit Parameters** and **Shares** tabs that present information concerning the various shares of the group. The **Information** tab presents the information for the whole group, while the **Shares** tab shows a list with detailed information for each share. Right-click on one of the column headings and select **Column properties...** if you wish to add or remove information columns from the table. Totals for the share group is shown at the top of the **Shares** tab:

Shares	Number of shares in the group
Archived files	Number of files that have been archived
Total files	Total number of files stored in the group's shares
Total volume	Total volume of the data stored in the group's shares
Fulfilled files	The number of files for which the requirements of the storage policy have been fulfilled
Not fulfilled files	The number of files for which the requirements of the storage policy have not been fulfilled.

From the **Shares** tab it is also possible to **Add**, **Move**, and **Delete** shares from the group.

You can use the group's **Edit Parameters** tab to define common **Parameters** for all the shares in the group. These parameters are explained under [“Share Information”](#) and [“Modify Share Parameters”](#).

Move a Share

It is possible to move shares from one Share Group to another and between the root of the shares tree and a share group.

To move a share, right-click on it and select **Move**. Select a destination and click **OK**.



Note

Share groups are only visual representations to aid the administrator in organizing shares. The share location in the file system or the user access to it is not affected.

To move a share from a Share Group to the root, check the **Move to root** box (the list of share groups will become greyed out) and click **OK**.

Rename a Share

You can change the name of a share. To do so, select the share in the **Shares** view and then select the **Edit Parameters** tab for the share. Type the new name in the **Name** field. To save the change, click on the **Apply** button at the bottom of the tab window.



Important

If the share you want to change is currently activated on one or more nodes, the share must be deactivated before the name change can take effect. When a share is deactivated, all user sessions on the share will be disconnected. A confirmation message will be displayed. Click on **Yes** to deactivate the share and apply the name change.

To make the share available again after the name change, you must activate the share as described under [“Activation on Nodes”](#).

Once the share is reactivated, users can connect to the share using the new name.

Delete a Share

If a share is no longer needed, it can be deleted. Right-click on the share and select **Delete**. A dialog box will display asking you to confirm the action. If you click on **Yes**, a warning message will then display to inform you that all files and folders on the share will be deleted. If you click on **Yes** again, the deletion process begins.

During the deletion process, the share is renamed by adding the current date to its name and it will remain visible in the **Shares View** until all the metadata and the file data relating to it are removed from the system.

Depending on the amount of data to be deleted, this process could take a long time.



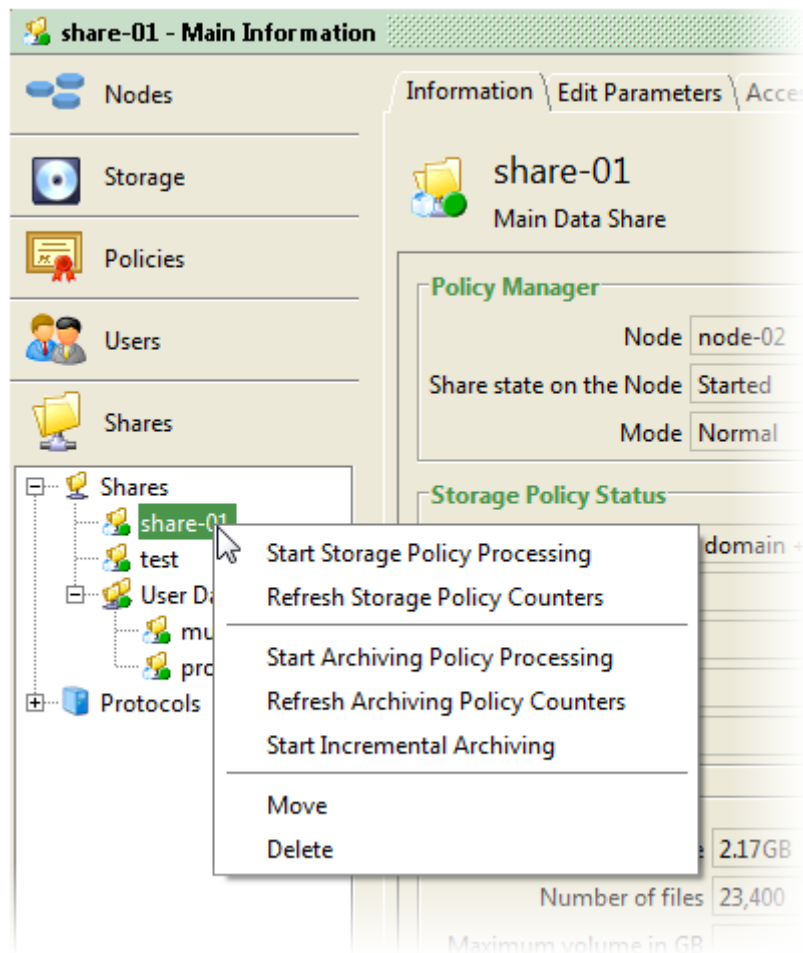
Warning

There is no *Trash bin* or *Recycling bin*. Once a share is deleted, its contents are *irrecoverable*.

Policy and Archive Processing

Additional options are available when you right-click on a share. Which options are available and which are greyed-out depends on which policies are associated with the share or whether and not the share is in a group. These options can be useful in the event of an error or for verifying that the policy is functioning correctly.

Figure 23.5. Shares Context Menu



Start Storage Policy Processing

This command overrides the scheduling set for the share and manually starts the processing of the storage policy.

Refresh Storage Policy Counters

This command resets to zero the counters displayed in the **Information** tab for a selected share, thereby forcing the system to recalculate their values.

Start Archiving Policy Processing

This command overrides the scheduling set for the share and manually starts the processing of the

	archiving policy. It checks if each archive has reached its retention duration, and if so marks it as "expired".
Refresh Archiving Policy Counters	This command resets to zero the counters displayed in the Information tab for a selected share, thereby forcing the system to recalculate their values.
Start Incremental Archiving	This will launch an automatic archiving process for those items in the share that have not yet been archived.
Move	Used for moving the share to or from a Share Group. See "Move a Share" for more information.
Delete	Used for deleting a share.



Warning

Note that if you use the **Delete** option on a share it will delete the share and all files and folders within that share. *Use with extreme caution.*

Policy Status Indicators

Visual status indicators help identify the location of any potential issues affecting the shares. Since a share can have both an archive policy and a storage policy at the same time, the indicators are separate for each type of policy. The archive policy status is indicated by small modifiers on the left side of the share icon, in the shape of a box. The storage policy modifiers are shown as colored dots on the right of the share icon.

Table 23.1. Policy Status Indicators

Policy Indicator		Color	Meaning
Archive	Storage		
		White (archive) Green (storage)	Policy requirements satisfied.
		Orange (blinking)	Policy is being processed or verified after a start-up of the node.
		Red	There is an error or the policy cannot be processed at this time (for example if the replication schedule is set to only allow processing after a certain point in time).

Share Information

Because shares are central to the functionality of the Active Circle system, the **Information** tab for an individual share contains detailed information that is divided into five areas:

- [Policy Manager](#)
- [Storage Policy Status](#)
- [Archiving Policy Status](#)
- [Quotas](#)

- [Versioning Volume](#)
- [Parameters](#)

Figure 23.6. Shares Information tab

The screenshot shows the 'share-03 - Main Information' window with the 'Information' tab selected. The left sidebar shows a tree view with 'share-03' highlighted. The main content area displays the following settings:

- Policy Manager:**
 - Node: node-01
 - Policy processing Node: <Automatic> ☒
 - Share state on the Node: Started
 - Mode: Restricted
- Storage Policy Status:**
 - Storage Policy: 1 copy
 - Fulfilled files: 0
 - Not fulfilled files: 0
 - New files: 0
 - Obsolete files: 0
- Archiving Policy Status:**
 - Archiving Policy: Archive-pol1
 - Archived files: 0
 - Not archived files: 0
- Quotas:**
 - Data volume: 0Bytes
 - Number of files: 0
 - Maximum volume in GB:
 - Maximum number of files:
- Versioning Volume:**
 - Volume: 0Bytes
 - Number of versions: 0
- Parameters:**
 - File access Audit: ☒
 - Enable data Replication Scheduling: ☒
 - Enable Distributed Lock Manager: ☒

Policy Manager

The **Policy Manager** area contains high-level information concerning the Policies associated with the selected share. The following fields are displayed:

Node

Node responsible for managing the Policies associated with the selected share. The Policy Management service (see [Chapter 27, Policy Management](#)) is responsible for managing which nodes are designated for processing the policy for the various shares.

Share state on the Node

The current status of the selected share. Possible values are **Started**, in which case the share is available to end users, and **Off**, in which case it is not.

Mode

The mode in which the node is currently operating. Possible values are **Normal**, in which case the data on the share is processed normally as indicated

Policy processing Node

in the policy rules; and **Restricted**, in which case only newly added files will be processed. For more information on the **Restricted** mode, please refer to [“Edit Parameters”](#) in [Chapter 27, Policy Management](#). The preferred node for processing the policy. If set to `<automatic>`, then the node chosen by the Policy Management service will be assigned the task.

Storage Policy Status

The **Storage Policy Status** contains the following fields:

Storage Policy	The name of the particular storage policy (if any) associated with the selected share.
Fulfilled files	The number of files (and versions of those files) for which the requirements of the storage policy have been met.
Not fulfilled files	The number of files (and versions of those files) for which the requirements of the storage policy have <i>not</i> been met. If this number is zero, the share has a green icon next to it.
New files	The number of files not yet processed by the policy.
Obsolete files	The number of files not subject to the policy (retention time expired, as defined by the storage policy).

Archiving Policy Status

The **Archiving Policy Status** contains the following fields:

Archiving Policy	The name of the particular Archiving Policy (if any) associated with the selected share.
Archived files	The number of files that have been archived by the Archiving Policy (if any) associated with the share.
Not archived files	The number of files that have <i>not</i> been archived by the Archiving Policy (if any) associated with the share.

Quotas

This section presents information concerning the **Quotas** for a selected share. To learn more about **Quotas** and how to configure them, please refer to [“Quotas”](#) later in this chapter.

Data volume	Quantity of data currently stored on the share.
Number of files	Number of files currently stored on the share.
Maximum volume in GB	Maximum amount of data allowed on the share as defined in the Quotas tab.
Maximum number of files	Maximum number of files allowed on the share as defined in the Quotas tab.

The [Refresh](#) button allows you to update the **Quotas** fields.

Versioning Volume

This section displays the total amount of data associated with the share and the total number of *file versions*. By way of an example, if there was only one file associated with a share, but there were 100 version of that file, then this field would display "100".


Parameters

For more information on these options, please see [“Parameters”](#) in the **Edit Parameters** tab.

Inherit parameter values

There are check boxes without text labels to the right of the value fields. These are *Inherit* check boxes. When one or more of these boxes are checked, it means the parameter will inherit the value for this field from its parent object (a group, a root object, etc.). You will find these boxes to the right of many fields throughout the interface, they are always used for inheriting values. When the *Inherit* box is checked, the original value field is greyed out. Some inherit fields are only displayed when a share belongs to a share group.

Figure 23.7. Inherit check boxes



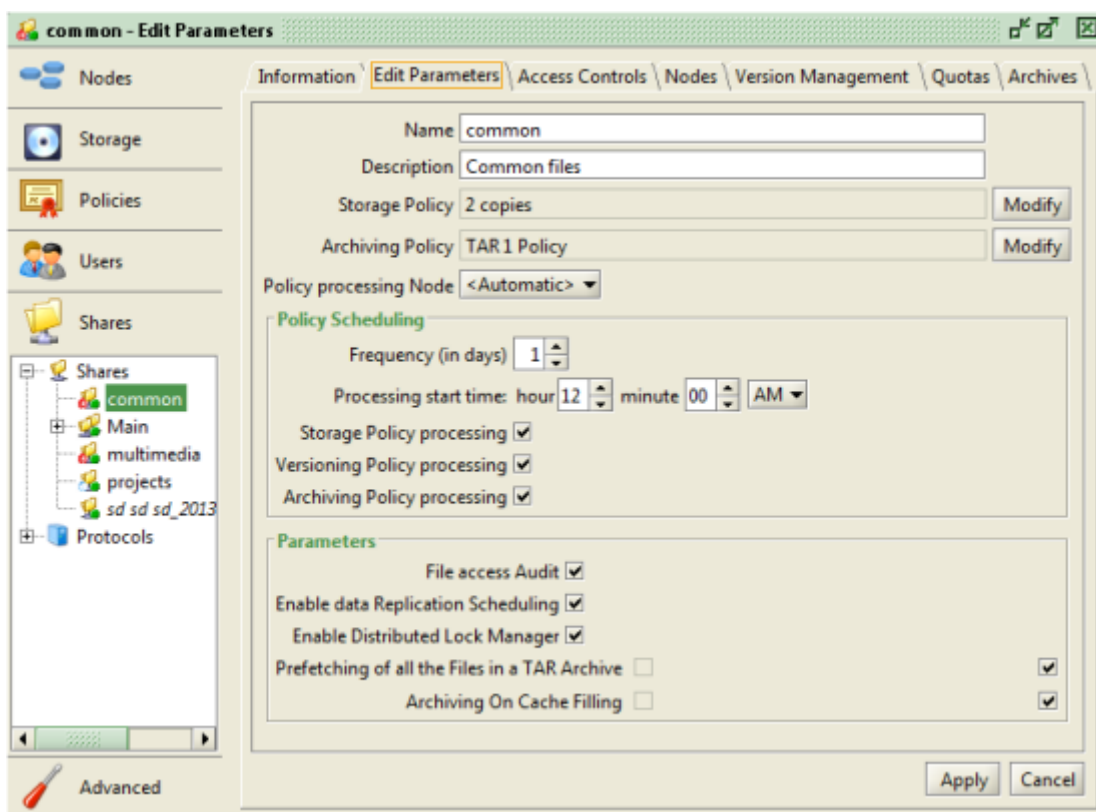
Modify Share Parameters

At the top of the **Edit Parameters** tab you can modify the **Name** and **Description** fields by entering the appropriate text in the fields.

The rest of the panel is divided into three sections:

- Options for [policy associations](#)
- [Policy Scheduling](#) options
- Audit, replication, lock manager, prefetching and automatic archiving [parameters](#)

Figure 23.8. Shares – Edit Parameters tab

**Tip**

Remember that any changes made must be validated by clicking the **Apply** button at the bottom of the tab.

Policy Association

You can add, modify, or remove the storage policy and/or the Archiving Policy associated with the share.

**Note**

Remember that a share must be associated with at least one policy.

To add or modify an association, click on the **Modify** button next to the policy, select the policy to associate, and click **OK**.

To remove an association, click on the **Modify** next to the policy to remove, check the box **No Policy**, and click **OK**.

Using the **Policy processing Node** drop-down, you can choose whether the processing of the Policies for the share should be assigned automatically or whether this task should be assigned to a particular node. If the preferred node happens to be off-line, then another node is chosen for the task.

Policy Scheduling

The **Policy Scheduling** area allows you to configure the time and frequency of the execution of the policy processing.



Important

In order for a type of policy to be automatically started by the policy scheduler, the check box for that type of policy must be ticked here. If no *Storage* or *Archiving* policy has been associated with the share, the check boxes for processing these policy types are greyed out.

Because policy processing can consume a considerable amount of system resources, it is recommended that you schedule these events for times when users are not connected to the system.

Furthermore, in systems with a large number of shares, it is recommended that you stagger the start-up of the policy processing in such a way that they do not all start at once on any given node.

Figure 23.9. Shares Policy Scheduling

Policies Scheduling

Frequency (in days)

Processing start time: hour minute

Storage Policy processing ☒

Versioning Policy processing ☒

Archiving Policy processing ☒

By default, policy processing is scheduled to run every day at midnight. To change the frequency, select a different number of days and click the **Apply** button. For example, changing the number to 2 will mean that policy processing takes place every other day.

To change the start time for the policy processing, adjust the **Hour** and **Minute** fields to reflect the time at which you wish the processing to begin. Click the **Apply** button to save your changes.

There are checkboxes for each type of policy processing. Checking the box next to a policy type activates the automatic start-up of the policy processing according to the schedule defined here.

Versioning Policy processing and **Storage Policy processing** are activated by default. If an archiving policy is associated with the share, the policy will not be run until you tick the box next to the **Archiving Policy processing** item in the **Policy Scheduling** area.



Note

If all the policy types are selected, the versioning policy is processed before the other policy types.

Parameters

Figure 23.10. Share Parameters

Parameter	Value	Inheritance
File access Audit	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Enable data Replication Scheduling	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Enable Distributed Lock Manager	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Prefetching of all the Files in a TAR Archive	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Archiving On Cache Filling	<input type="checkbox"/>	<input checked="" type="checkbox"/>

These additional check boxes provide more options for the share. You can select these parameters using two sets of check boxes. The box immediately following the text field is valid for this share only, while selecting the check box to the far right means this share will *inherit* this value from its parent object (some inherit boxes are only available when the share is a member of a share group).

File access audit

Activating this option will start a logging service that records information on file transactions such as date of creation, modification, deletion, who performed the action, etc. For information on accessing the data collected by the audit, please see [Chapter 25, Auditing and Statistics](#).

Enable data Replication Scheduling

Checking this box activates the replication schedule (see “[Replication Schedule](#)”) for inter-domain data transfers for the current share.

Enable Distributed Lock Manager

Checking this box activates the **Distributed Lock Manager** (see [Chapter 26, Distributed Lock Manager](#)) for the current share.

Prefetching of all the files in a TAR archive

This parameter activates the prefetch functionality for this share, which is off by default. Upon activation, when a file in a TAR archive on tape is requested, all the files in the TAR file will be read into the cache at once. This is a specific feature developed to address situations where an application continually reads all the files contained in a TAR file without being capable of sorting the read. This functionality is only available if the functionality that allows TAR files to be created according to folder depth in a tree structure is activated on the share (“structured TAR”). To activate the checkbox, you must first uncheck the inheritance checkbox on the right side of the window. You can use the associated circle property to activate this functionality for all shares:

Circle property: `vfs.prefetch`

Default value: false (Off)

This parameter only has effect if the “structured TAR” feature is enabled, either in the [advanced archive policy parameters](#)

Archiving on cache filling

or system-wide by setting the circle property `archive.closeTAROnFolderChangeDepth` to another value than "0" (default).

This parameter activates the functionality for automatic triggering of an archive operation when the cache reaches a certain filling level. This feature makes it possible to avoid cache saturation during large and continuous file deposits, as files will be automatically archived as the cache fills up. This functionality is only available for shares which are not associated with a storage policy, since storage policies automatically copy files from the cache to disk. The functionality is meant for shares which only have an archive policy. To activate the checkbox, you must first uncheck the inheritance checkbox on the right side of the window. You can use the associated circle property to activate this functionality for all shares:

Circle property:

`archive.triggerOnCacheFilling`

Default value: false (Off)

The cache filling limit can be defined by the number of files in the cache waiting to be archived or the data volume in the cache. If you need to change the default thresholds for this feature, use these circle properties:

Circle property:

`archive.triggerOnCacheFillingNumberThreshold`

Default value: 1,000,000

Circle property:

`archive.triggerOnCacheFillingVolumeThreshold`

Default value: 500 GB



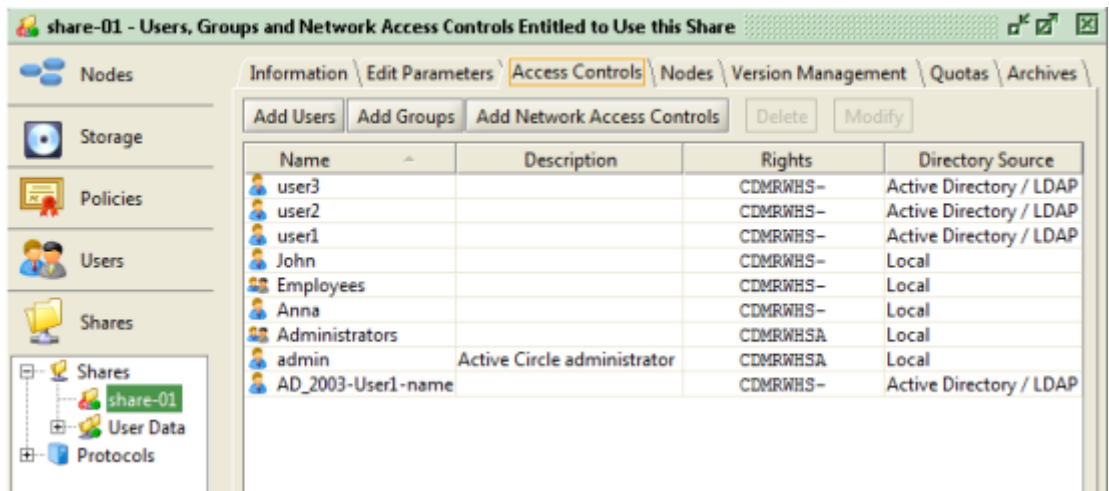
Tip

For information about how to modify circle properties, see [“Circle Properties”](#) in [Chapter 5, The Circle](#).

Access Controls

The **Access Controls** tab for a share is used for granting users, groups, and applications access to that share. The processes for performing these actions are detailed in the following sections.

Figure 23.11. Access Controls tab



For information on creating users and groups, and on how to use external directory services, please see [Chapter 18, User Profiles and Access](#).

Add Users, Groups, and Network Access Controls

The process for giving access to a share is the same, whether it be a user, a group, or a machine.

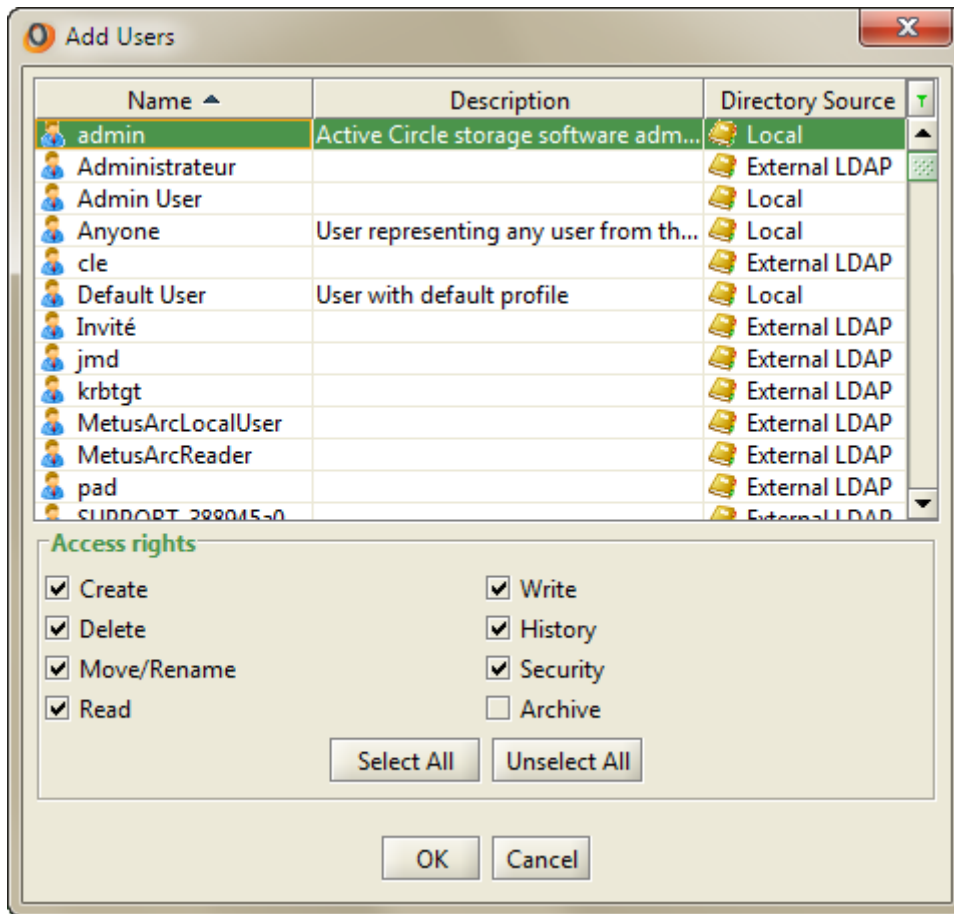


Note

The default user `admin` is automatically added to the Access Controls list when the share is created.

To provide access to a share, click on the button corresponding to the object to which you wish to grant access — [Add User](#), [Add Group](#), or [Add Network Access Control](#) — and then select the element (or elements) to which to grant access.

Figure 23.12. Add Users



Next you must decide which **Access Rights** to give the element. The options are:

Create	Write
Delete	History
Move/Rename	Security
Read	Archive



Note

All of the access rights are selected by default with the exception of **Archive**, which grants users the right to generate on-demand archives. For the built-in "admin" user, even the **Archive** right is selected by default.



Tip

The **Access Controls** tab is only used for controlling access to the share and cannot be used to control file-level security. For more information on managing file security in a Microsoft Windows environment, please refer to [Appendix D, Managing Windows ACLs](#).

Activation on Nodes

In order for a share to function it must be activated on one or more nodes. This can be done through the **Nodes** tab for the share.



Note

When a share is created, the system will automatically activate it on all the nodes of the circle.


To deactivate a share on a node, select the node in the list and click the **Deactivate Share** button. You can also reactivate a share, by clicking the **Activate Share** button.

Figure 23.13. Nodes tab

Name	Active	State	Policy Management
node-01	Yes	In Service	No
node-02	Yes	In Service	Yes
vcentos	No	In Service	No

If the file server of the node you activate the share on is not started, the system will ask you to confirm if you want to activate the file server now:

Figure 23.14. File Server Activation



If the node's file server is not running, the share will not be accessible through that node. For more information, see [“Activating the File Server”](#) in [Chapter 24, Protocols](#).



Tip

A share can also be activated on a node on the **Shares** tab for a node: **Shares View > Protocols > Domain > Node > Shares tab**.

For the share to function, the **Status** should read "Started" and the **In Service** column should read "Yes". Once started, the **Active** column will read "Yes".

Possible statuses for a node are:

- Started
- Off
- Ready to start
- Starting



Note

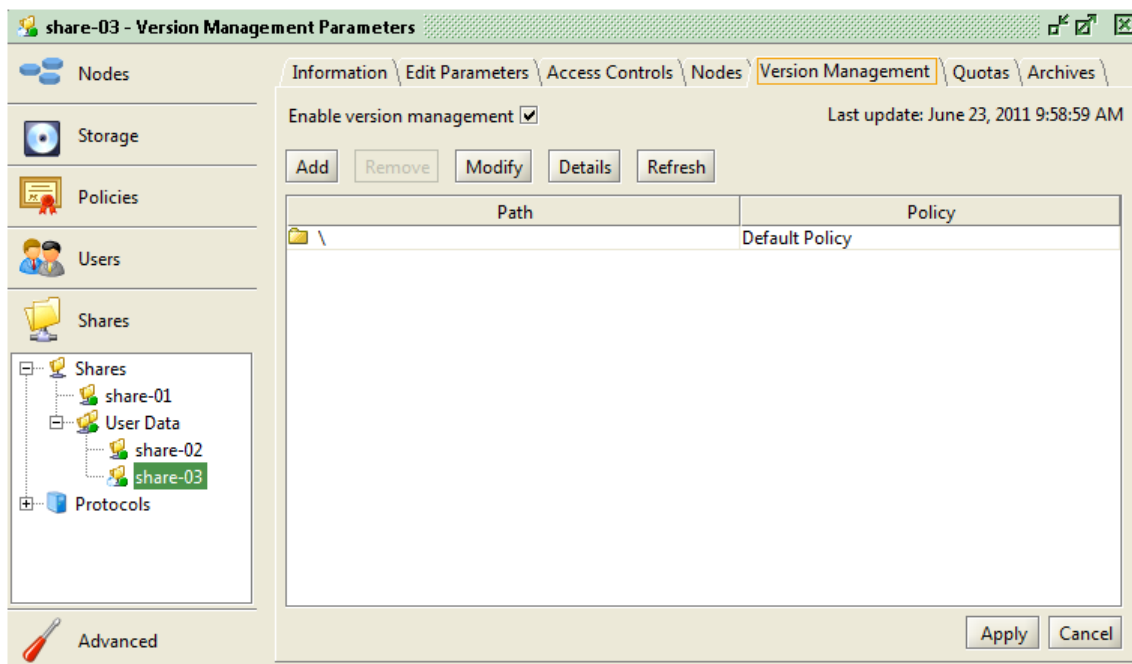
The status **Off** can indicate a problem on the node.

The **Policy Management** column indicates whether a given node has been assigned the task of managing the policies for the share.

Version Management

Version management is enabled by default when you create a share. You can deactivate it by unchecking the **Enable version management** checkbox.


Figure 23.15. Version Management tab



A table displays the folders in the share and the Versioning Policy associated with the individual folders. By default, only the root folder will appear with whichever Versioning Policy (if any) was associated with the share in the **Edit Parameters** tab.

You can add sub-folders of the share to the table in order to apply a different Versioning Policy to that folder and any of its sub-folders by clicking on the **Add** button.

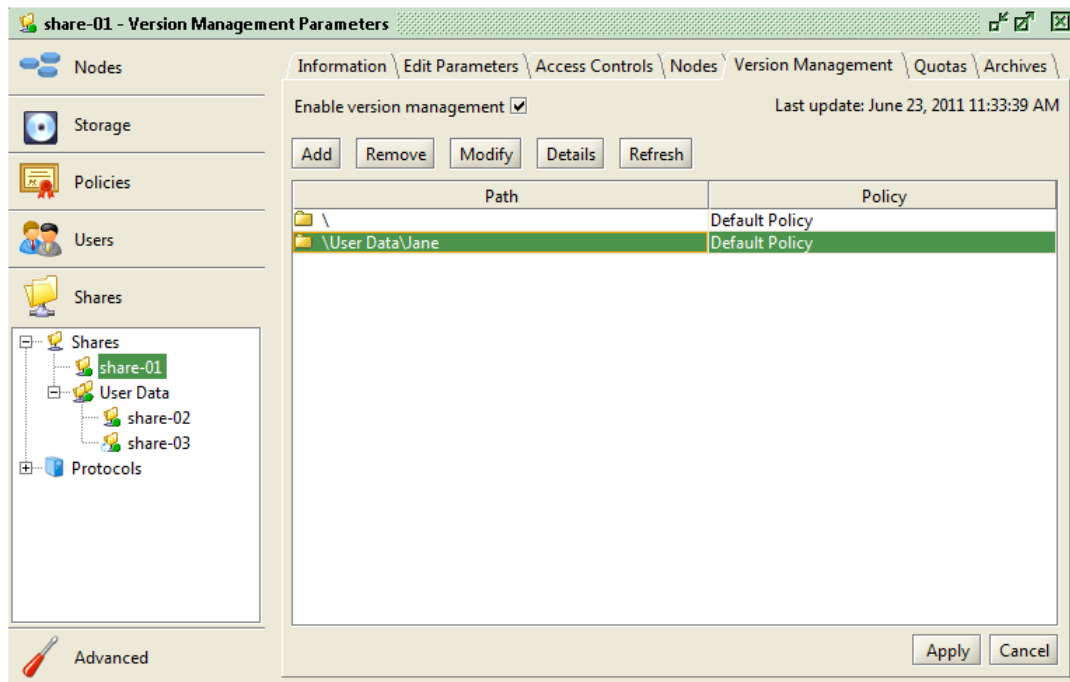
Figure 23.16. Add Sub-Folder



You can browse to the specific sub-folder to add to the table and then select the desired Versioning Policy. Click **OK** to close the dialog box.

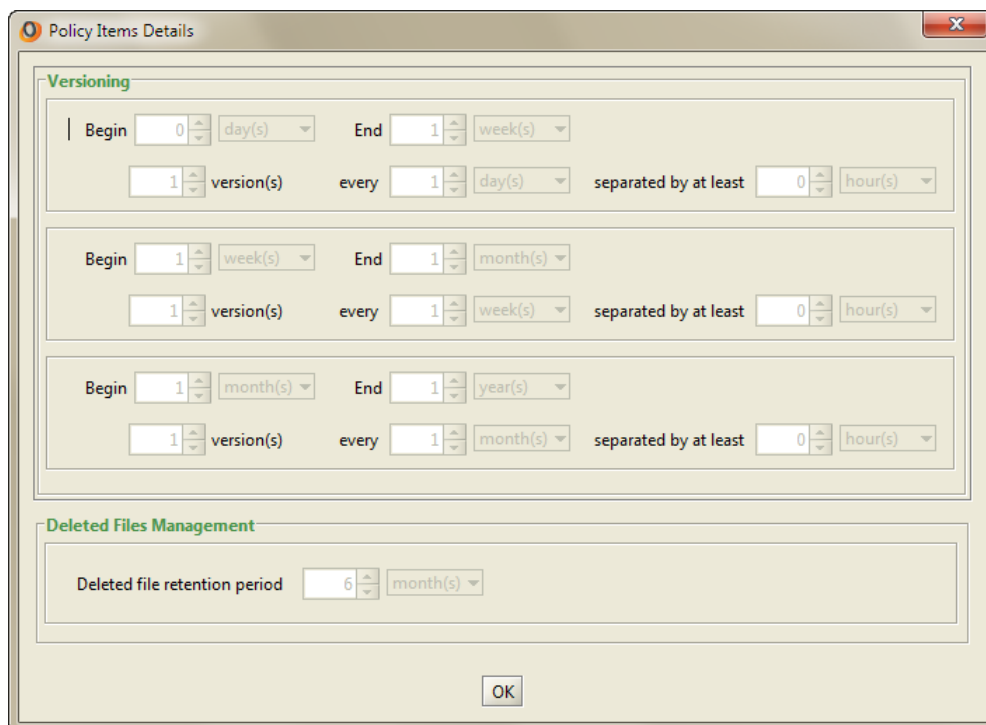
You can also change which Versioning Policy applies to a particular folder by selecting it in the table and clicking the **Modify** button, selecting the new Versioning Policy, and clicking **OK**.

Figure 23.17. Sub-Folder Added



Clicking on the **Details** button for a particular folder brings up a dialog box that lists the rules that have been defined for the selected Versioning Policy.

Figure 23.18. Versioning Policy Details

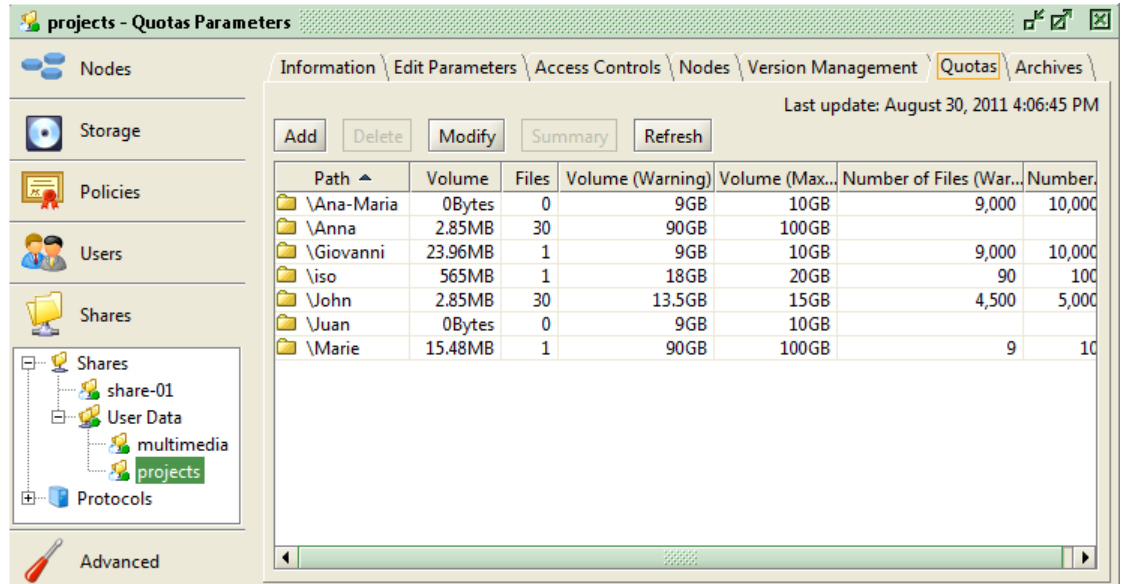


Quotas

You can use the **Quotas** feature for shares to provide a system for monitoring and controlling the use of storage space across your Circle. For example, by defining a maximum quantity of data or a maximum number of files allowed for a particular share or group of shares (or folders


within shares), an administrator can prevent individual users in an organization from taking up too much of the storage space in the system. You can set a warning level above which the system will send a **Supervision Note** to the administrator's in-box so that corrective action can be taken before the maximum volume is reached.

Figure 23.19. Quotas tab



The default size of a new share is set to 1 petabyte by use of a default quota on the root of the share. The size can be changed by modifying the quota set on the share root. This quota cannot be deleted.

Figure 23.20. Default Quota for Share Root



Important

If you set a maximum quantity of data or a maximum number of files, this is considered an absolute limit and if it is reached, the system will not allow users to create new versions or to add new files and a **Supervision Note** (with a level of Major) is sent to the administrator.

Rules for Setting Quotas

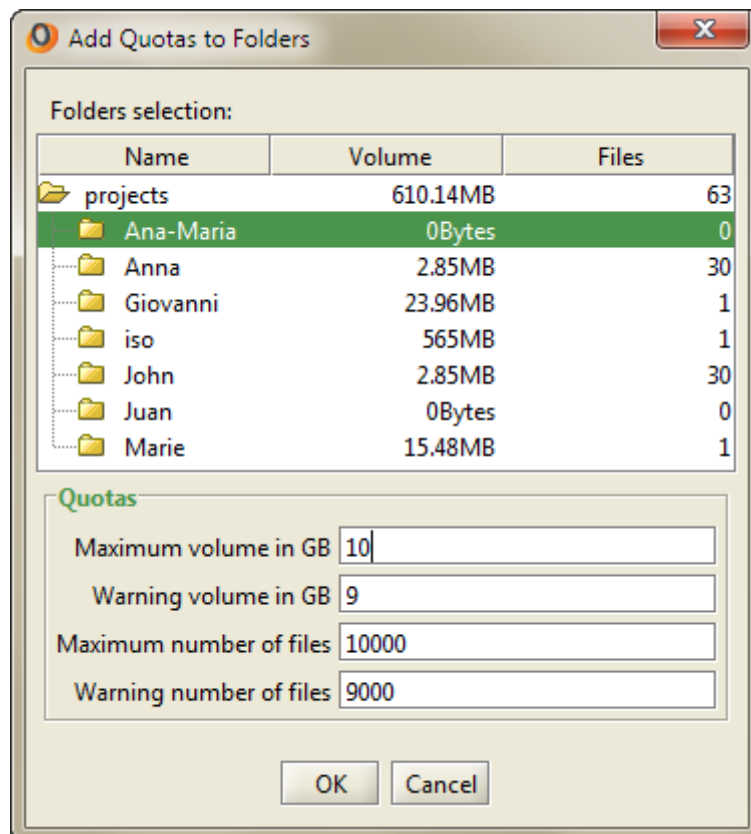
- You do not have to define values for both **Maximum Data Volume** and **Maximum Number of Files**. You may define either, both, or none at all.
- If you define a maximum value for either data volume or number of files, you must also define a warning level. This value is automatically set by default to 90% of the value entered into the maximum value field, but you can change it.

- The analysis of the contents of a folder take into account its contents and that of all of its sub-folders.
- When a limit is set for a particular folder, you may not set the maximum value for a sub-folder to a value higher than its parent.
- The calculation is based on the most recent versions of the files in the folder and sub-folders (the versions visible to the users). Deleted files and the version management history are not taken into account.
- The **Quotas** set for a folder remain associated with it even if the folder is moved or renamed.
- The values for the folders are not recalculated immediately following an upgrade of the Active Circle software. They will be recalculated the next time the Versioning Policy is executed for the share.

Setting Quotas

To set a **Quota** for a share or folders, click on the **Add** button. A dialog box appears that presents a navigable tree structure with the share as its root. You can set the **Quota** by selecting either the share, in which case the **Quota** will apply to it and all of its sub-folders, or you can select a sub-folder.

Figure 23.21. Add Quotas to Folders



Once you have selected a folder, you must then set the values for the **Quota**. You can set quotas for two parameter types: **Data Volume** and **Number of Files**. You must enter values for both the **Maximum** and the **Warning** fields if you want to specify quotas.

The available fields are:

- **Maximum data volume (in GB)**

- **Warning data volume (in GB)**
- **Maximum number of files**
- **Warning number of files**

Once you have entered values for one or both parameters, click **OK** to set the **Quota**. The dialog box will close and the information concerning the **Quota** will appear in the table.

Archives

The **Archives** tab shows a list of all the archives and allows you to monitor the state of the archives and archive processing for a given share. You can see a list showing the same information for all the shares in the system by selecting the **Archives** tab for the root element **Shares**. The list of archives is by default sorted in reverse chronological order by the **Creation Date** (newest on top). Click on a column header to sort on another field. The configuration of the Archiving Policy is covered in [Chapter 17, Archiving Policy](#).

Figure 23.22. Archives tab

Name	State	Phase	Status	Type	Creation Date	Duration
Archive:multimedia#4	Successful	-	Complete	Incremental	August 26, ...	0:00:03:58
Archive:multimedia#1	Interrupted on error	Waiting for tape	Empty	On demand	August 24, ...	0:00:00:02
Archive:multimedia#2	Successful	-	Complete	On demand	August 24, ...	0:00:00:09
Archive:multimedia#3	Successful	-	Complete	On demand	August 24, ...	0:00:01:16

The top part of the tab contains four fields providing information about the overall state of the archives for the current share:

Archives	The number of archives that have been started for the share (also counts incomplete and interrupted archives).
Successful	The number of archives in the state <i>Successful</i> .
Archived files	The total number of files that have been archived.
Volume	The total volume of the archive data for this share.

A row of buttons provides additional controls over the archive selected in the table.


Interrupt	Stops the selected archiving process. Only available if the process is running.
Delete	Deletes the archive. <i>Please note that there is no warning message.</i>
Location Details	Detailed information on the physical location of the tape on which the archive is stored (if available).
Files	The files contained in the archive.
Summary	The Summary provides the following information for the selected archive: number of archives, number of successful archives, number of archived files, and the volume of the archive.

RefreshRefreshes the data displayed in the **Archives** table.

Location Details

To see information on the location of an archive, click on the **Location Details** button. A window is displayed, showing which pool, tape(s) and TAR file(s) the archive is located in:


Figure 23.23. Archive Location Details



Files in the Archive

To see the list of the files included in an archive, click on the **Files** button. A window is displayed, showing a tree view of all the files in the archive, including the file version:

Figure 23.24. Archive File List



Click in the **Only files not archived** check box to display only the files which were excluded from the archive.

Archive Details

The table in the **Archives** tab provides detailed information about the archives for the current share.

Column	Description
Name	Name of the archive
State	The current state of the archiving process. Possible values are: <ul style="list-style-type: none"> • Not scheduled • Scheduled • In progress • Cancelled • Interrupted on demand / Interrupted on error / Interrupted on service stop • Completed with errors • Successful • Time quota exceeded
Phase	The different phases of the archiving operation. In order, they are: <ul style="list-style-type: none"> • Building list • Waiting for tape • Writing • Recording • " - " (when the archive is finished).
Status	Status of the archive. Possible values are: <ul style="list-style-type: none"> • Complete • In progress • Valid • Expired (as defined by the retention period defined in the Archiving Policy). • Empty (when interrupted or error)
Type	Indicates if the archiving type was Incremental, On demand or Export
Creation Date	Date the archive was created
Duration	Length of the archiving process
Volume	Size of the archive
Archived Files	Number of files in the archive (during an archiving in progress, it shows the number of files archived so far, in real time)
Not Archived Files	Number of files excluded from the archive process (when an archiving is in progress, this column shows the number of files not yet archived, in real time)
Location	The name of the <i>Tape Pool</i> where the archive is located.

When an archiving process begins, the **Status** is **In progress** and the **Phase** passes from **Building list** to " - " when the processing is finished. The **State** should be **Successful** and the **Status** should read **Complete**. During an archiving process, you can monitor the progress in real time in the columns **Archived Files**, **Not Archived Files** and **Volume**.

If for some reason the archiving process is interrupted, the reason for the interruption is listed in the **State** column and the **Phase** column retains whatever value it had when the process was interrupted.

Archive Display Filter

You can control which types of archives to display in the table, using the row of check boxes below the table to the right:

Figure 23.25. Archive Filter check boxes

Automatic ☒ Import ☒ Export ☒ On demand ☒

By default, all the boxes are checked. When there are a large number of archives in the table, it can be useful to remove the display of archive types you are not interested in. For example, if you want to only see the **On demand** archives, you remove the check mark in the other three check boxes by clicking in them.

Manual Archiving

By using the Archiving Policy, you can ensure that periodic archives are automatically created. It is, however, also possible to start an archiving process manually. To initiate a manual archiving process, right-click on a share in the **Shares View** and select **Start Incremental Archiving**.



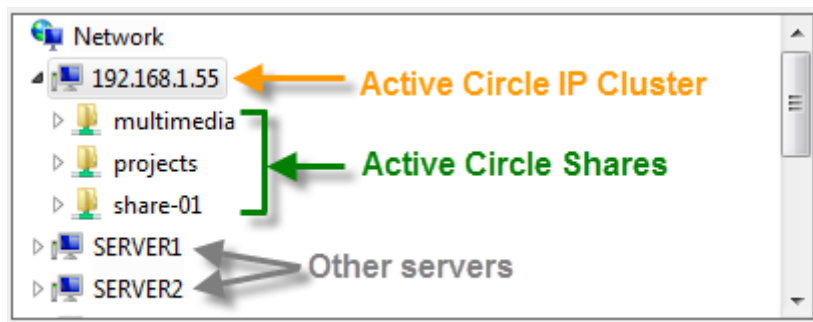
Note

If for some reason a file that is scheduled to be archived is not, a Supervision note is sent to the administrator as notification of the problem.

Connecting to Shares

Unlike a traditional Windows share or Linux/Unix mount point, an Active Circle share will not be visible as a folder in the directory structure on the server hosting the node. Because an Active Circle share is a component of the Virtual File System, it is only visible to an end user via a *mapped drive* or a *mounted folder* on the client machine. Users log on to the Active Circle shares from their clients just like to any other network share, using the CIFS (Windows), NFS or FTP protocols. An example of each connection type is shown below.

Figure 23.26. Example – Active Circle Shares available through Network Drive Mapping on a Windows Client (CIFS)



Example – FTP connection to the IP cluster "192.168.1.55":

```
$ ftp 192.168.1.55
Connected to 192.168.1.55 (192.168.1.55).
220 FTP server ready
Name (192.168.1.55:user): Kurt
331 User name okay, need password for Kurt
Password:
230 User logged in, proceed
ftp> ls
drwxrwxrwx  1 -1 -1 0 Jan  1 1970 multimedia
drwxrwxrwx  1 -1 -1 0 Jan  1 1970 projects
ftp> ls multimedia
-rwxrwxrwx  1 0 0 70326 Sep 30 11:38 Procedures.pdf
-rwxrwxrwx  1 0 0 8871  Oct 18 14:20 Project Template.odt
drwxrwxrwx  1 0 0 0      Sep 26 17:42 Videos
drwxrwxrwx  1 0 0 0      Sep 26 17:43 Images
```

With FTP, the user logs directly on to a node or a cluster. All the shares the user is authorized to access and that are activated on the specified node (or cluster), will be available to the user (in this case the shares "multimedia" and "projects" as listed by the **ls** command).

Example – NFS connection (mount) from a Linux client to the share "multimedia":

```
$ mount 192.168.1.55:/multimedia /mnt/multimedia
$ ls -l /mnt/multimedia/
drwxrwxrwx 1 user user  512 Sep 26 17:43 Images
drwxrwxrwx 1 user user  512 Sep 26 17:42 Videos
-rwxrwxrwx 1 user user 33792 Oct 18 14:24 Thumbs.db
-rwxrwxrwx 1 user user  8871 Oct 18 14:20 Project Template.odt
-rwxrwxrwx 1 user user 70326 Sep 30 11:38 Procedures.pdf
```

In the above example, for the user issuing the **ls** command (or any other command/operation on the share) to be granted access, the UID of that user must be added to the access controls of the share (type **id** on the Linux command line to display the UID and GID of the current user). This means that a user with this UID must exist in the Active Circle directory and that this user (or a group containing this user) must be in the list of users in the [Access Controls tab](#) for the share. By default, the **root** user on Linux does not have access to NFS shares in Active Circle (this setting can be overridden by disabling the field **Enable root_squash** in the NFS parameters, see the "NFS" section in [Chapter 24, Protocols](#)).

Recommended NFS/CIFS connection options on GNU/Linux:

When connecting to Active Circle shares from a GNU/Linux client, certain mount options are recommended. Note that there should be no line breaks in the command lines, those are only used for readability and space purposes below.

- **CIFS:**

```
mount -t cifs //serverName/shareName
-o username=userName,password=userPassword
,file_mode=0777,dir_mode=0777,mapchars,nounix,nouser_xattr
```

- **NFS:**

```
mount node:/share mountPoint -o hard,intr,acdirmin=1,tcp,vers=3
```



Note

In addition to access rights, in order for users to be able to access an Active Circle share, two basic conditions must be met:

- The share must be activated on at least one node in the Circle
- The share must be associated with one storage policy and/or one Archiving Policy

Depositing or Migrating Files to Shares

When the shares and the file servers have been configured and activated, the shares can be populated with user files. Users who connect to the shares can deposit files using regular file copy or drag and drop in the operating system or save files to the shares from the applications they use. The files deposited will automatically be subject to the replication, versioning and archiving rules defined in the system. No administrator intervention is needed for this. It may, however, be useful for an administrator to copy files from existing shares/directories or external media to populate Active Circle shares. Two common ways of doing this are:

- Migrating files using Robocopy (see [Appendix E, Migrating Files to Shares](#))
- Importing a file structure from tape (see ["Importing an Archive"](#))

Chapter 24. Protocols

The **Protocols** section in the Shares view presents the view of the nodes in your Circle similar to that found in the Nodes view, showing the domains, sub-domains and nodes in a tree structure. However, whereas the Nodes view is where the node storage properties are configured, the **Protocols** section under the Shares view is where you configure its networking properties.

The server is created automatically when a node is created. When first viewing a newly created node in the Shares View, its icon appears greyed out because the file server has not been activated. It is a storage device without a means to add or access files in the system. A Circle must have at least one file server activated.

Here you can specify which nodes in the Circle will provide users access to the system by hosting the shares (as opposed to servers operating simply as storage nodes and/or participating in policy processing).

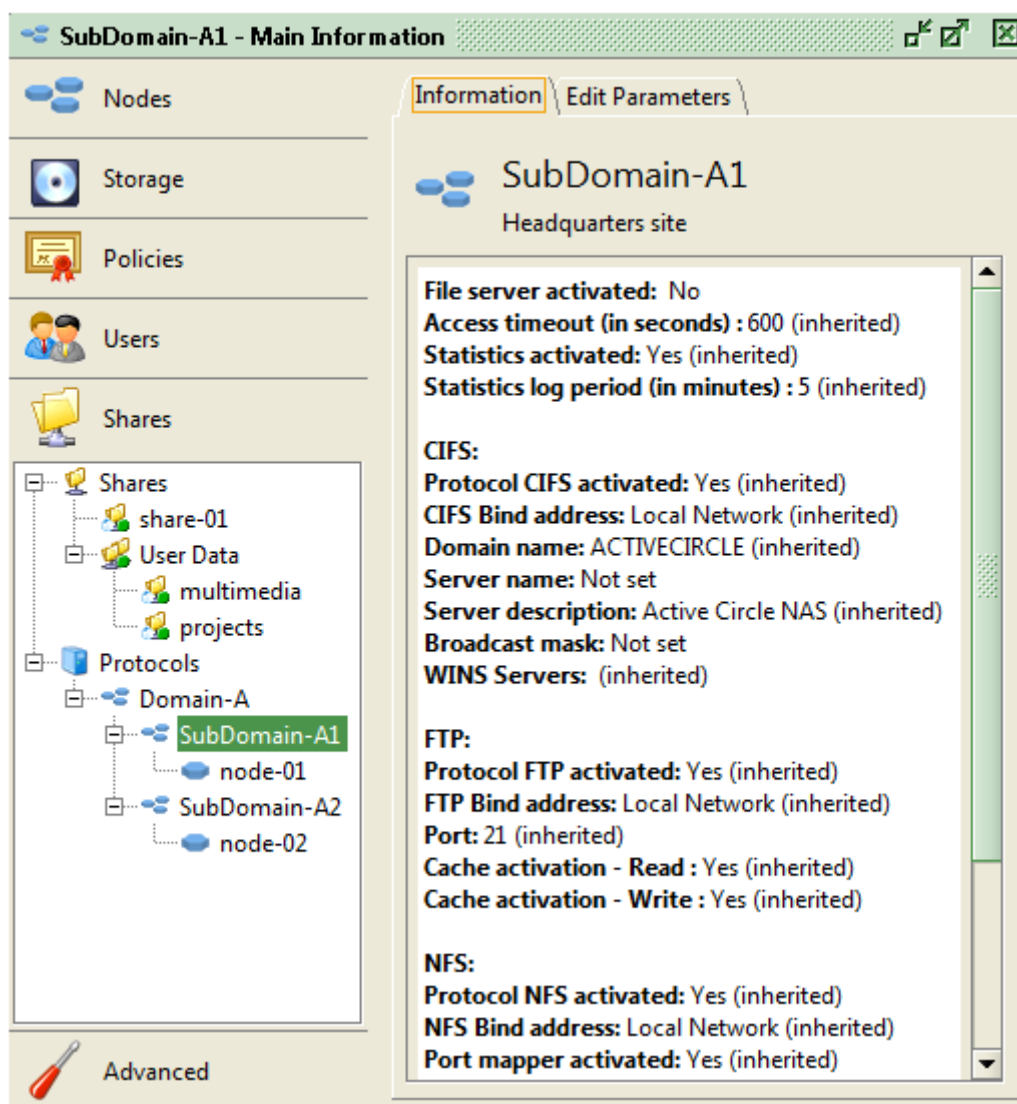
Additionally, you can configure the protocols used for accessing the files:

- [CIFS](#) (Windows)
- [FTP](#)
- [NFS](#)

In Active Circle, a *file server* means an Active Circle node with the file server service activated. By default, file servers inherit their configurations from the root file server element. Consequently, if you are administering a homogeneous system in which all users are accessing files in the same way, you can configure file server access at the top-most level and the configuration will apply to all of the nodes in the tree structure under the Circle.

If, on the other hand, you are managing a more diverse environment in which users and/or applications access the files using different protocols, then you must either configure each node individually or by domain or sub-domain.

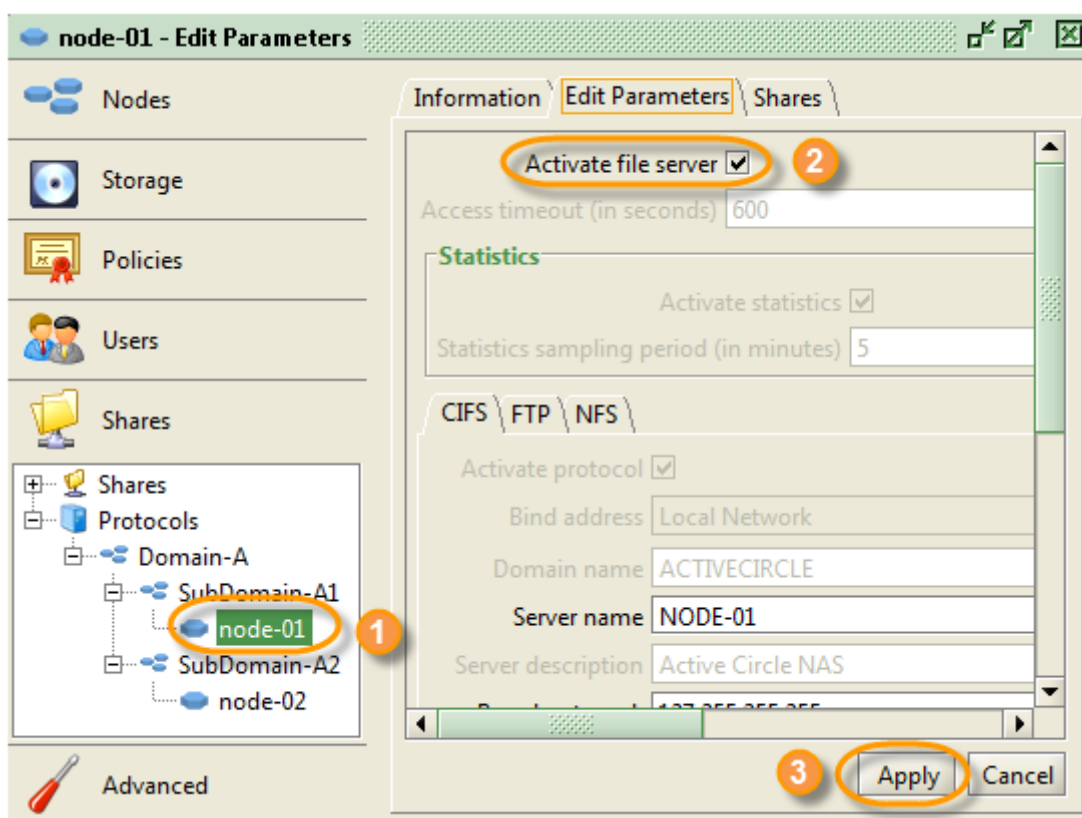
Figure 24.1. File server protocols



Activating the File Server

The system automatically activates the file server when a node is created. If a file server has been deactivated, it can be activated by expanding the tree structure and selecting the node that you would like to activate. Then click on the **Edit Parameters** tab. At the top of the tab there is a check box labelled **Activate file server**. Check the box and click on **Apply**. The node will appear as active under **Protocols**.

Figure 24.2. Activate File Server



When you activate a file server you automatically enable the three supported network file system protocols — CIFS, FTP and NFS — on that server. You may disable one or two of the protocols, depending on how the end users will access the files in the system.

When the file server has been activated, users can connect to it through the defined shares and start depositing and using files on it (assuming the shares and their corresponding user access rights have been properly set up). The next section describes how to configure the protocols for connecting to the file server, if the default settings are not appropriate for your environment.



Note

Any time you activate a share on a node where the file server is not started, you will be prompted by the system to activate the file server at the same time. If you answered 'Yes' to that question, the file server for the node will already be activated. For more information, see the section on [share activation](#) below or "[Activation on Nodes](#)" in [Chapter 23, Shares](#).

Restarting the File Server

If you need to restart the file server, this is done by first deactivating the file server and then activating it again on the **Edit Parameters** tab as described above. Follow these steps:

1. Remove the check mark from the **Activate file server** box by clicking in it.
2. Click on the **Apply** button at the bottom of the window. This stops the file server.

3. Start the file server again by clicking the **Activate file server** box so that a check mark appears in it. To activate the change, click on the **Apply** button at the bottom of the window.

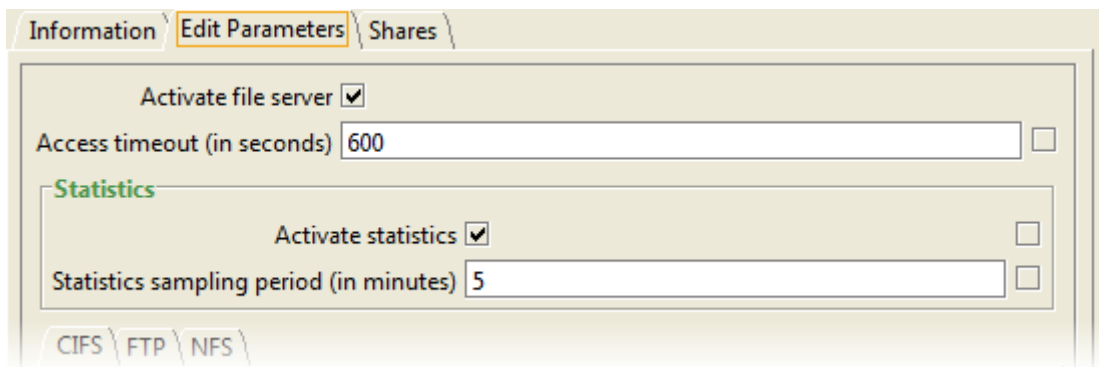
Configuring the File Server

Use the **Edit Parameters** tab for a node to configure the file server that is running on the node. The panel is divided into two sections:

- Settings for **Access Timeout** and **Statistics** at the top of the tab.
- Tabs for configuring each of the network protocols **CIFS**, **FTP** and **NFS** at the bottom.

You may configure file servers either individually or collectively. By default, each file server inherits its configuration from its parent element, meaning that the input fields in the **Edit Parameters** tab are greyed out. To modify the configuration for an individual file server, uncheck the box to the right of the parameter you wish to modify (this activates the line so that you can make changes).

Figure 24.3. Edit Parameters tab – General Parameters



If you make any changes, click the **Apply** button for the changes to take effect.

Access Timeout

Use this field to specify the maximum time in seconds allowed for read, write or seek requests on a file before the system sends an error to the client. The default value is 10 minutes (600 seconds). If the timeout is reached, a supervision note will be sent to the administrator. This would normally only occur with offline files on tape, heavy stress situations on a tape library, tape drive malfunctions or (if the timeout has been set very low) with extremely large files on disk.

Statistics

Use the fields under **Statistics** to activate and modify the settings for statistics logging.

Activate statistics

Checking this box activates the statistics logging feature for the file server. The logs generated contain information relating to the load on the server in terms of number of requests for different types of data for the statistics sampling period.

For detailed information concerning the content of the statistics logs, please see [Chapter 25, Auditing and Statistics](#).

Statistics sampling period	Length of time in minutes between two statistics measurements. Five minutes by default.
-----------------------------------	---

CIFS

Across the bottom there are three sub-tabs, one for each of the supported network file system protocols.


The **CIFS** (Common Internet File System, also called SMB or Server Message Block) protocol is commonly used for providing users working in a Microsoft Windows environment access to a file system over a network. Note that in the image below, all of the inherit checkboxes (the column of boxes to the right of the entry fields) have been unchecked to make the information easier to read in this document (in a typical configuration, some of the checkboxes would be selected). Note that when you deselect an inherit checkbox, an asterisk (*) is placed to the left of the text label for the corresponding field to indicate that there has been a change which has not yet been applied. By default, several of the protocol parameters are inherited from the root **Protocols** item.

Figure 24.4. Edit Parameters tab – CIFS

Activate protocol	Enables the protocol for the currently selected file server.
Bind address	Associates the service with the specified network address. The value by default is Local Network , meaning the file server will accept CIFS connections on all available network interfaces. If you want to bind the service to a particular address, uncheck the inherit checkbox and click Select . You may bind the service to the network interface, in which case it will use whatever is the default IP address for that interface, or directly to the IP address itself. See Figure 24.5, "Bind Address Selection" .
Domain name	NetBIOS name of the domain or workgroup to which the file server belongs.

Server name	The NetBIOS name by which users will be able to identify the file server on the network. By default, this is the node name, but you can change it. Maximum of 15 characters, no spaces.
Server description	A description to help identify the file server
Broadcast mask	Mask used in order for the node to be recognized in the network.
WINS servers	Servers providing WINS (Windows Internet Name Service). Blank by default. Optional.

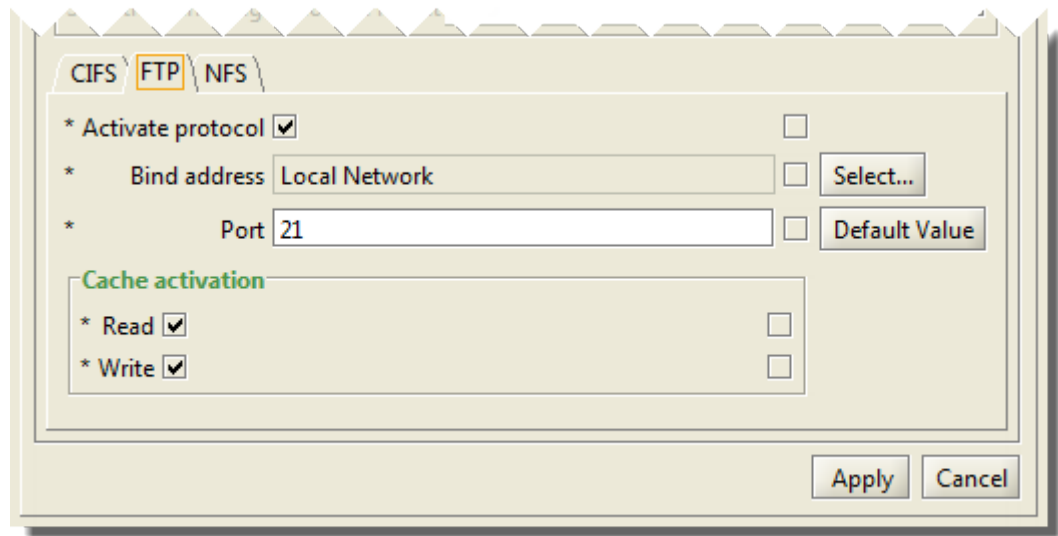
Figure 24.5. Bind Address Selection



FTP

In this tab you can configure the options for the **FTP** (File Transfer Protocol) service. Like the other protocols, it is activated by default at the root level. Note that in the image below, all of the inherit checkboxes (the column of boxes to the right of the entry fields) have been unchecked to make the information easier to read in this document (in a real configuration, all or some the inherit checkboxes might instead be checked). Note that when you deselect an inherit checkbox, an asterisk is placed next to the text label for the corresponding field to indicate that there has been a change which has not yet been applied. By default, several of the protocol parameters are inherited from the root **Protocols** item.

Figure 24.6. Edit Parameters tab – FTP

**Activate protocol**

Enables the protocol for the currently selected file server.

Bind address

Associates the service with the specified network address. For more information, please refer to the item of the same name in “CIFS” above.

Port

FTP port number. You can change this if required by unchecking the inherit checkbox and entering the desired value and then clicking the **Apply** button. If you want to return the port to its default value, click the **Default Value** button and then click **Apply**.

Cache Activation

The **Cache Activation** feature allows you to control whether or not files that are being written to a tape pool using the FTP protocol are first written to the Active Circle Cache.

Under normal circumstances, when a user puts a file on an Active Circle share, it is first written to the cache directory. The system regularly scans the cache for new files and when detected, they are processed into “packs” and copied to a storage pool where they are stored in Active Circle safes. The process also works in the other direction when reading files. When a user requests a file, Active Circle first checks if the file is still in the Cache. If it is found, it is served to the user from there. If not, the system retrieves the file from the pool, unpacks it to the cache, and then serves up the file.

The advantage of this is that when packing files into safes, Active Circle can optimize the storage by packing multiple files into individual safes. If on the other hand you are handling primarily very large files that will each occupy multiple safes, it might be more efficient to bypass the cache and write the files directly to tape, although they are still packed into safes.

The cache is activated by default for both read and write modes. By deactivating either or both, the system will ignore the cache and perform the operation directly to or from tape.

NFS

The **NFS** tab allows you to configure the behavior of the Network File System protocol. Note that in the image below, all of the inherit checkboxes (the column of boxes to the right of the

entry fields) have been unchecked to make the information easier to read in this document (in a real configuration, all or some of the inherit checkboxes might instead be checked). Note that when you deselect an inherit checkbox, an asterisk is placed next to the text label for the corresponding field to indicate that there has been a change which has not yet been applied. By default, several of the protocol parameters are inherited from the root **Protocols** item.

Figure 24.7. Edit Parameters tab – NFS

The screenshot shows the 'Edit Parameters' window for the NFS protocol. The 'NFS' tab is active. The window contains the following settings:

- Activate protocol:** Checked (checkbox).
- Bind address:** Local Network (text field).
- Enable port mapper:** Checked (checkbox).
- Port mapper port:** 111 (text field).
- Mount server port:** 35032 (text field).
- Lock server port (NLM):** 35033 (text field).
- Server port:** 2049 (text field).
- Nfsd count:** 16 (text field).
- Encoding:** UTF-8 (dropdown menu).
- Manage GIDs:** Unchecked (checkbox).
- Enable root_squash:** Checked (checkbox).

Below these settings is a table for 'Root squash exception (no_root_squash)' with columns 'Name' and 'Description'. There is an 'Add' button and a 'Delete' button below the table. At the bottom right are 'Apply' and 'Cancel' buttons.

Activate protocol

Enables the protocol for the currently selected file server.

Bind address

Associates the service with the specified network address. For more information, please refer to the item of the same name in “CIFS” above.

Enable port mapper

Checkbox for enabling the port mapper service. Activated by default. This service should be deactivated if the port mapper service of the operating system is running and you prefer to use that instead. If the port mapper of the operating system is already running and you activate the Active Circle port mapper, Active Circle will detect it and use the port mapper of the operating system.

Port mapper port


Port used by the port mapper service. Port number 111 is used by default.

Mount server port

Port used by the Mount server. Port number 35032 is used by default.

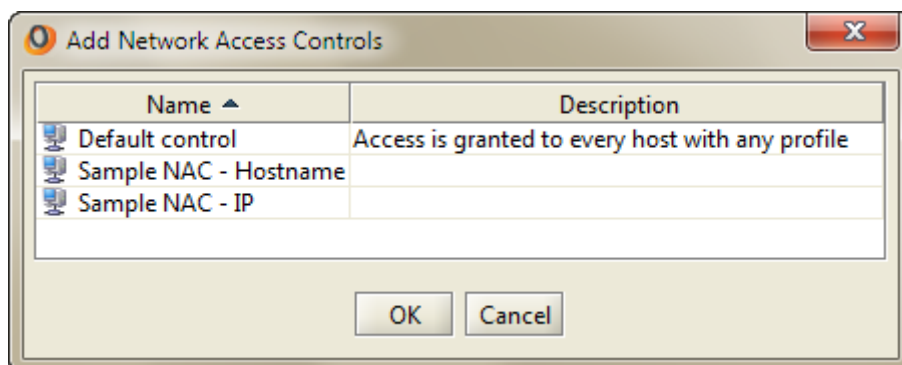
Lock server port (NLM)

Port used by the NLM (Network Lock Manager) service. Port number 35033 is used by default.

Server port	Port used by the NFS server. Port number 2049 is used by default.
Nfsd count	The number of nfsd threads launched by the NFS daemon. Note: For client machines running IRIX (from SGI), this value must be set to 1 for those servers on which the NFS will be mounted.
Encoding	Character encoding used for file names. By default UTF-8 (Unicode) is used. ISO-8859-1 (Latin alphabet with accents) is also available.
Manage GIDs	This option can be used to bypass the limit of 16 auxiliary group ids (GIDs) imposed by the NFS protocol. NFS requests normally contain a user id and a primary group id along with a list of auxiliary group ids. All these user credentials are used to authorize access. Due to a limitation in the NFS protocol, at most 16 auxiliary group ids can be sent by the client. If you check the Manage GIDs box, the list of auxiliary group ids received from the client will be replaced by the list of auxiliary groups provided by the imported external directory for authentication purposes. This is therefore a way to provide complete authorization for users who are members of more than 16 groups. Note that the primary group id is not affected so a new group command on the client will still be effective. If you change this setting, the new value will only be taken into account by new NFS sessions. To create new NFS sessions, all existing NFS sessions on the node(s) must first be closed using the Active Circle Administration Tool (for more information, see “Sessions” in Chapter 7, Nodes). The next access by each NFS client will automatically and transparently create a new session with the updated value. This option is disabled by default.
Enable root_squash	<code>root_squash</code> is enabled by default in order to prevent root users on client machines from having privileged access to exported files. It is possible to disable this (that is, to give root users privileged access) or to create exceptions for <code>root_squash</code> (see below).
Root squash exception (no_root_squash)	<p>This option allows you to create exceptions to the <code>root_squash</code> rule while maintaining the security of having <code>root_squash</code> enabled.</p> <p>To create a root squash exception, uncheck the inherit checkbox to enable the feature and then click the  Add button to select a previously defined Network Access Control (see Chapter 22, Network Access Controls for more information). Select the client host controls for which you would like to enable root access and click OK. The selected machine(s)</p>

will appear in the list of machines allowed the `root_squash_exception`.

Figure 24.8. Add Network Access Controls

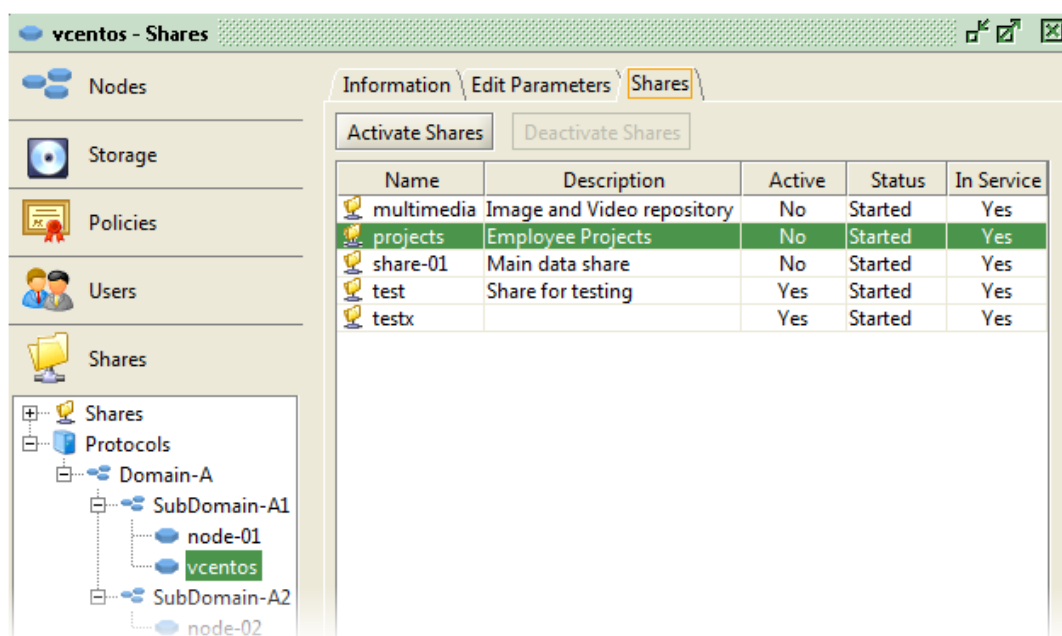


Activating Shares

Once the file server has been activated and configured, you must also activate one or more shares on the file server in order for users to be able to access and create resources there. This can be done either through the **Shares** tab for a selected file server or through the **Nodes** tab for a selected share (see [“Activation on Nodes”](#) in [Chapter 23, Shares](#)).

To activate a share, select it in the list and click the **Activate** button. The entry in the **Active** column for the share will switch to **Yes**. The **Status** and **In Service** columns should read **Started** and **Yes** respectively.

Figure 24.9. Activate a Share



Chapter 25. Auditing and Statistics

This chapter describes the *Share Auditing* and *Server Statistics* features and the log files generated by them. Each service is activated separately. When activated, the system logs information about file access and/or server activity on each node.

Share Auditing

A service that logs file and directory events such as date of creation, modification or deletion, changes in size, attributes, mode or UID/GID, and file read or write. For each event, information about user, protocol, IP address, path and changed data is recorded.

Server Statistics

A feature that collects information on server activity in terms of the number of requests for certain types of data for the indicated period of time.

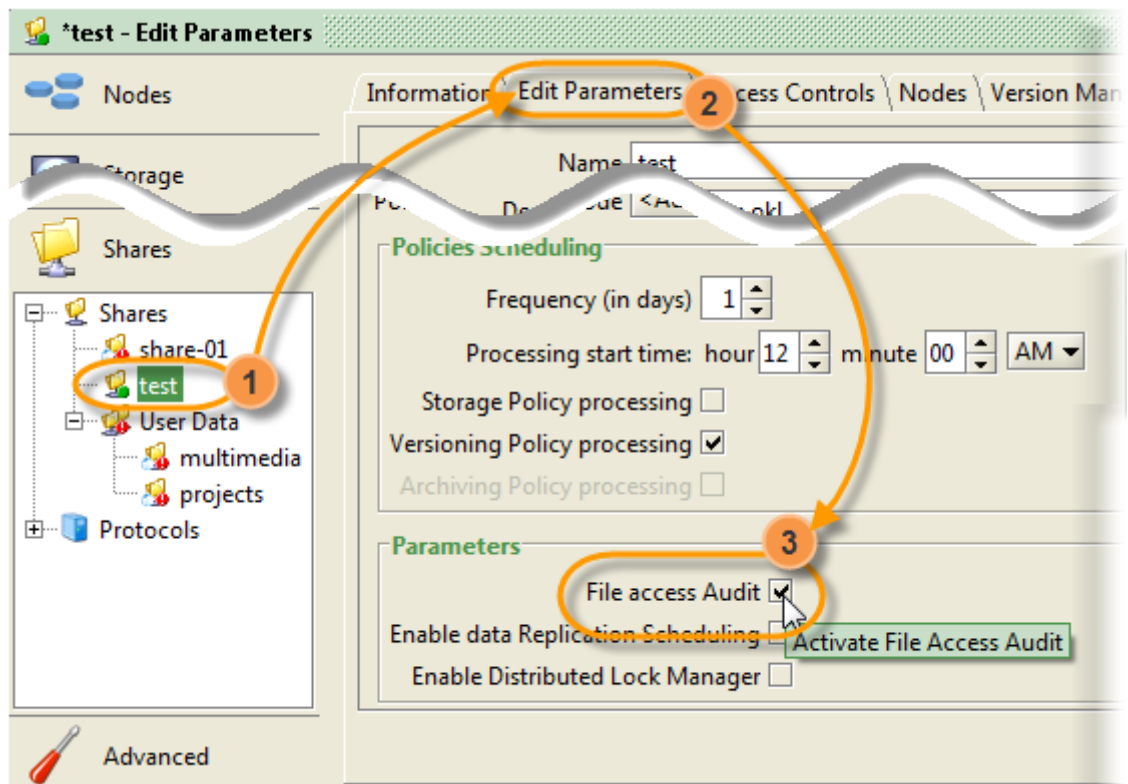
Both logging types are based on the same functionality and will therefore be described together in this chapter. Where the functionality differs (such as activation and certain log details), this will be indicated in the text or described in separate sections.

Activation

Share Audit Activation

The share auditing functionality is activated by default. It can be deactivated or activated individually for each share in the Active Circle Administration Tool using the option **File Access Audit** in the **Parameters** section of the **Edit Parameters** tab of a share.

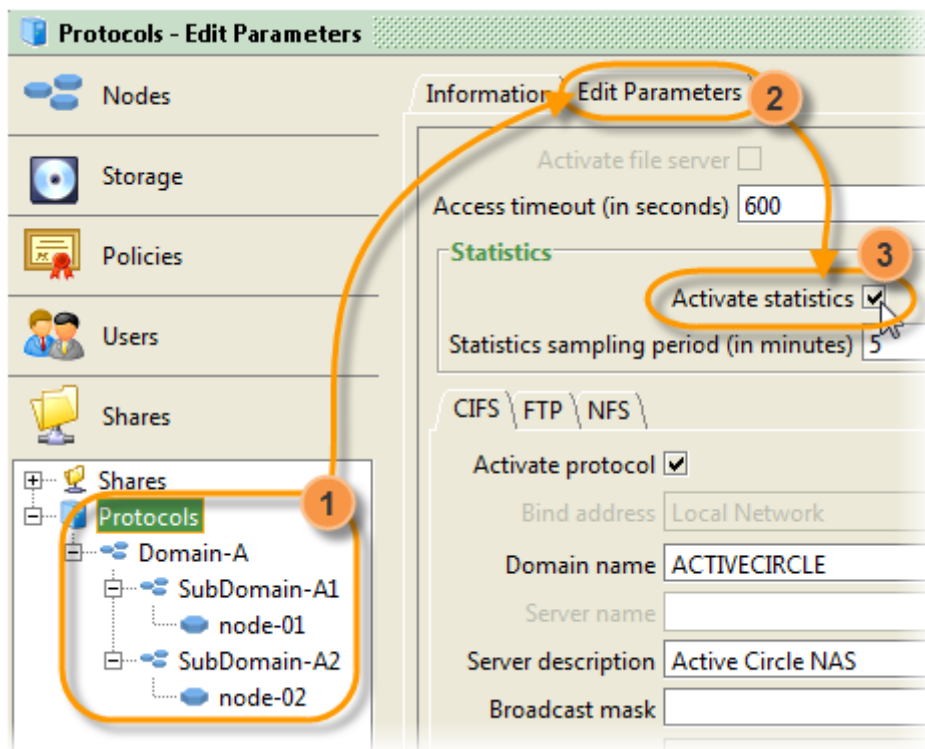
Figure 25.1. Share Audit Activation



Server Statistics Activation

Server statistics collection is activated at the level of the node, subdomain, domain or the circle. By default, this functionality is activated and inherited from the parent object. In the Active Circle Administration Tool, go to the **Shares** view. Expand the tree under **Protocols**. Click on any of the items in the tree, depending on where you want to activate the functionality. Then select the **Edit Parameters** tab in the right pane, and tick the check box for the option **Activate statistics**. For example, if you select the top-level item **Protocols**, the activation will apply to all the nodes in all the domains of the circle (the default), unless you have disabled inheritance of this option at any level of the circle (this is done by unchecking the box to the far right of the **Activate statistics** field). If you select a specific node or domain, the activation will only apply that one node or domain.

Figure 25.2. Server Statistics Activation



Just below the activation field, you can modify the sampling interval for the statistics log, by entering a new number of minutes in the **Statistics sampling period** field. Each line in the log file will contain a summary of the activity during the period of this interval (5 minutes by default).

Log File Characteristics

The log files generated by the **Share Auditing** and **Server Statistics** are in comma separated value (CSV) format and can be opened in a text editor or any spreadsheet program such as OpenOffice Calc or Microsoft Excel.

The log files are organized by share and by day, with a new file being generated every day at midnight (this process is called "rollover"). The previous day's log file is automatically compressed using the *gzip* format. In addition, Active Circle has a system for rolling over the files when they grow too large.

**Note**

An **audit log** is created only when there is activity on a share. If a whole day passes without file or folder activity on a share, there will be no audit log for that date. The statistics logs, however, are created on each node every day, recording summary information even for sampling periods when there has been no activity.

**Tip**

This section contains several references to *Circle Properties*. These are properties which can be used to set and modify default values at the node, domain or circle level. For more information on how to modify these properties, see “[Circle Properties](#)” in [Chapter 5, The Circle](#).

Location

Statistics log files are managed by share. A separate set of logs is produced and stored for each share. For audit logs, a [consolidated mode](#) is enabled by default (does not apply to statistics).

Default log file location:

```
/activecircle/cell/data/Stats/NAS/shares/share-name
```

Log file location in consolidated mode (audit only):

```
/activecircle/cell/data/Stats/NAS/shares/
```

File naming convention (where *yyyy-mm-dd* is the date):

Audit:	<code>nas-audit_YYYY-mm-dd.csv.gz</code>
Statistics:	<code>nas-usage_YYYY-mm-dd.csv.gz</code>

The log files are handled in a similar way to the administration log messages. The logs are specific to each node and there is no automatic consolidation of the log files for a Circle.

**Note**

Activity is logged on the node where the user connects to the share. If the share access happens directly on a node, the logs for the user activity will be written to that node. It is therefore possible to have logs with the same date (but with different content) on several nodes. If the share is accessed through a cluster (see [Chapter 28, HA Cluster Services](#)), the log entries will be written to the leader node of the cluster.

Size limits

There is a size limit for each log file and for all the log files combined. When a log file reaches its size limit (or when the date changes), it is compressed by default using the gzip format and a new log file is opened.

If the *combined size limit* for all the files is reached (for the share and the auditing or statistics functionality), logging is suspended. A [supervision note](#) with warning level *Minor* is generated as notification about the suspension. During the suspension of logging, message collection

continues in memory. When the amount of space occupied by the log files goes back under the limit as a result of the hourly purge, the logging starts again automatically and any data collected during the suspension is written to file. You can modify the size limits and the compression setting using *circle properties*:

Log file size limits and compression properties

Default size limit:	10 MB
Size limit property:	<code>nas.activityFileMaxLength</code>
Default combined size limit:	100 MB
Combined size limit property:	<code>nas.activityFileMaxAggregateLength</code>
Compression property:	<code>nas.activityFileCompression</code>
Name of compressed gzip file:	<code>nas-audit_YYYY-mm-dd[_x].csv.gz</code> where the x in <code>[_x]</code> is a positive integer indicating the version of the file, if more than one file is created in a day due to the file size limit being reached. The file without the <code>[_x]</code> will be the most recent.



Note

For the **audit log**, a safeguard has been put in place to limit the number of events accumulated in memory in these potential situations:

- when logging is suspended
- when events are generated faster than the auditing service is able to write them to the log file

The default value for the maximum number of audit events in memory is 40,000. **Beyond this number, events are lost forever and will not be logged to file.** A supervision note is issued as a warning if this limit is being approached.

The memory limit can be modified using this property:

`nas.audit.maxPendingEvents`

Duration and logging interval

The files are purged from the system when they reach the limit of the defined file retention period. The age of the files is verified when the functionality is activated, when a log file rollover is performed, and every hour. The age limit can be changed using a circle property.

Default file retention limit:	100 days
Retention limit property:	<code>nas.activityFileRetentionTime</code>

The current **audit log** file is updated either every 5 seconds or, if there are more than 100 events within the preceding second, every second.

The **statistics counters** are collected for the duration of the sampling period (five minutes by default) after which they are written to file and the counters reset to zero. To modify the sampling period, see [Statistics activation](#).

The updating interval for the audit and statistics logs is not verified if the file is locked by a third-party application (such as Microsoft Excel). In this case, events and statistics will be collected in memory and written to the log after the application has released the file.

Consolidation (audit only)

A consolidated mode is enabled by default for events of all shares on a node. The audit log files will be created at the node level, and one combined set of log files will be written to the location `/activecircle/cell/data/Stats/NAS/shares`. This mode does not apply to server statistics.

Circle property for enabling audit consolidation: `nas.audit.consolidated`

In consolidated mode, the combined file size limit that will result in suspension of logging is calculated by multiplying the value of the combined size limit property (`nas.activityFileMaxAggregateLength`) with the number of shares in service (including shares where the audit feature is not enabled). Logging will be suspended if this value is exceeded.



Note

The consolidated mode is automatically enabled when operation sessions are in monitored mode. This is managed by the circle property `session.monitored`. If this is set to `true`, FTP sessions will be monitored and the audit log service will enter the consolidated mode.

Log File Contents

Audit Log

Each line in the audit file corresponds to an event. For each event, the following data is collected:

DATE	The date in ISO 8601 format.
TIME	The time in ISO 8601 format.
time_t(ms)	The number of milliseconds that have passed since 1 January 1970.
NODE (CELL)	The name of the node where the data was collected.
SHARE	The name of the share where the data was collected.
PROTO	The name of the protocol used to generate the event. Possible values:
	FTP If the share is accessed via FTP and operation sessions are being monitored (the circle property <code>session.monitored</code> is set to <code>true</code>).
	- In all other cases.
EVENT	The event identifier:
	F.CREA / D.CREA New file/directory created
	F.WRIT The file was modified at least once before being closed (this id was previously called F.CHGE). This event is only generated when the file is closed, and only if it

	was modified. Unlike for other events, the DATE and TIME correspond to the closing of the file. This event is also generated in direct I/O mode.																		
F.READ	The file was read at least once before being closed. This event is only generated when the file is closed, and only if it was read. Unlike for other events, the DATE and TIME correspond to the closing of the file. This event is also generated in direct I/O mode.																		
F.DELE / D.DELE	File/directory deleted																		
F.SIZE	File size changed																		
F.ATTR / D.ATTR	File/directory attributes changed																		
F.MODE / D.MODE	File/directory mode changed																		
F.UID / D.UID	File/directory UID changed																		
F.GID / D.GID	File/directory GID changed																		
F.RENA / D.RENA	File/directory renamed																		
F.MOVE / D.MOVE	File/directory moved																		
USER	The name of the user who performed the operation. If the system cannot determine the user, a dash "-" is displayed.																		
FROM	The IP address of the client machine where the operation was executed. If the system cannot determine the IP address, a dash "-" is displayed.																		
PATH	The path to the file or directory (presented in UNIX/Linux format). In order to comply with the CSV standard, the names of the files are presented in quotes ("/foldername/filename"). If the file name itself contains quotes, then they are doubled.																		
NEW	Information in the NEW column depends on the type of EVENT (where F=file and D=directory): <table> <tr> <td>F.RENA / D.RENA</td><td>The new name after renaming</td></tr> <tr> <td>F.MOVE / D.MOVE</td><td>The new name after a move</td></tr> <tr> <td>F.SIZE</td><td>The new file size</td></tr> <tr> <td>F.ATTR / D.ATTR</td><td>The new attribute value</td></tr> <tr> <td>F.MODE / D.MODE</td><td>The new mode value in octal</td></tr> <tr> <td>F.UID / D.UID</td><td>The new UID (in decimal)</td></tr> <tr> <td>F.GID / D.GID</td><td>The new GID (in decimal)</td></tr> <tr> <td>F.WRIT</td><td>The total number of bytes written between the opening and the closing of the file</td></tr> <tr> <td>F.READ</td><td>The total number of bytes read between the opening and the closing of the file (the value might be null if the read attempts failed). In direct I/O mode, the number of bytes is estimated using the size of the file(s) being transferred. Consequently, in direct I/O mode the value can be higher (but never lower) than the actual number of bytes read.</td></tr> </table>	F.RENA / D.RENA	The new name after renaming	F.MOVE / D.MOVE	The new name after a move	F.SIZE	The new file size	F.ATTR / D.ATTR	The new attribute value	F.MODE / D.MODE	The new mode value in octal	F.UID / D.UID	The new UID (in decimal)	F.GID / D.GID	The new GID (in decimal)	F.WRIT	The total number of bytes written between the opening and the closing of the file	F.READ	The total number of bytes read between the opening and the closing of the file (the value might be null if the read attempts failed). In direct I/O mode, the number of bytes is estimated using the size of the file(s) being transferred. Consequently, in direct I/O mode the value can be higher (but never lower) than the actual number of bytes read.
F.RENA / D.RENA	The new name after renaming																		
F.MOVE / D.MOVE	The new name after a move																		
F.SIZE	The new file size																		
F.ATTR / D.ATTR	The new attribute value																		
F.MODE / D.MODE	The new mode value in octal																		
F.UID / D.UID	The new UID (in decimal)																		
F.GID / D.GID	The new GID (in decimal)																		
F.WRIT	The total number of bytes written between the opening and the closing of the file																		
F.READ	The total number of bytes read between the opening and the closing of the file (the value might be null if the read attempts failed). In direct I/O mode, the number of bytes is estimated using the size of the file(s) being transferred. Consequently, in direct I/O mode the value can be higher (but never lower) than the actual number of bytes read.																		
Other cases	-																		

Chapter 25. Auditing and Statistics

OLD Lists the values before the change (with the exception of the actions F.RENA, D.RENA, F.MOVE and D.MOVE).

Example 25.1. Share Audit File (CSV format)

```
DATE,TIME,time_t(ms),CELL,SHARE,PROTO,EVENT,USER,FROM,PATH,NEW,OLD
2011-09-26,17:35:33,131...964,node-01,share-01,FTP,D.CREA,admin,192.168..., "/Folde...
2011-09-26,17:35:55,131...698,node-01,share-01,-,D.RENA,admin,192.168..., "/Folder"...
2011-09-26,17:36:08,131...194,node-01,share-01,-,F.CREA,admin,192.168..., "/Misc/te...
2011-09-26,17:36:08,131...439,node-01,share-01,-,F.SIZE,admin,192.168..., "/Misc/te...
...
2011-10-12,11:51:11,131...915,node-01,share-01,-,F.READ,admin,192.168..., "/bkp/fil...
2011-10-12,11:51:12,131...356,node-01,share-01,-,F.SIZE,admin,192.168..., "/bkp/fil...
2011-10-12,11:51:47,131...148,node-01,share-01,-,F.WRIT,admin,192.168..., "/bkp/fil...
```

(Three dots means characters or lines have been removed for presentation purposes.)

Figure 25.3. Share Audit CSV file imported in a spreadsheet

	A	B	C	D	E	F	G	H	I	J	K	L
1	DATE	TIME	time_t(ms)	CELL	SHARE	PROTO	EVENT	USER	FROM	PATH	NEW	OLD
2	2011-09-26	17:35:33	1317051333964	node-01	share-01	FTP	D.CREA	admin	192.168.1.50	/Folder	-	-
3	2011-09-26	17:35:55	1317051355698	node-01	share-01	-	D.RENA	admin	192.168.1.50	/Folder	/Misc	-
4	2011-09-26	17:36:08	1317051368194	node-01	share-01	-	F.CREA	admin	192.168.1.50	/Misc/test06.txt	-	-
5	2011-09-26	17:36:08	1317051368439	node-01	share-01	-	F.SIZE	admin	192.168.1.50	/Misc/test06.txt	8910	0
6
7	2011-10-12	11:51:11	1318413071915	node-01	share-01	-	F.READ	admin	192.168.1.50	/bkp/CReg.reg	56	-
8	2011-10-12	11:51:12	1318413072356	node-01	share-01	-	F.SIZE	admin	192.168.1.50	/bkp/CReg.reg	140792974	0
9	2011-10-12	11:51:47	1318413107148	node-01	share-01	-	F.WRIT	admin	192.168.1.50	/bkp/CReg.reg	140823158	-



Note

During periods of high file activity on a share, a large number of audit events may be logged. For example, F.READ or F.WRIT will be logged each time a file is closed, when data are read or written in the file. This can lead to the log filling up and logging being suspended sooner than anticipated. To avoid this, you can deactivate one or more audit log events. Use the circle property `nas.audit.filteredOutEvents` to specify the events to be filtered out of the log. The values should be specified as a hexadecimal bitmask, where each event has the following value:

1=F.CREA	20=F.GID	400=D.ATTR	8000=D.RENA
2=F.WRIT	40=F.DELE	800=D.MODE	10000=F.READ
4=F.ATTR	80=F.RENA	1000=D.UID	20000=F.MOVE
8=F.MODE	100=D.CREA	2000=D.GID	40000=D.MOVE
10=F.UID	200=D.CHGE	4000=D.DELE	80000=F.SIZE

Examples:

- To filter out F.WRIT and F.MODE, specify "A" (hex addition 2+8=A)
- To filter out F.CREAT, F.WRIT, D.MOVE and F.SIZE, specify "C0003" (hex addition 00001+00002+40000+80000=C0003)

Statistics Log

Each line in the statistics file corresponds to data collected during a sampling period (which is 5 minutes by default). For each period, the following data is collected:

Chapter 25. Auditing and Statistics

DATE	The date in ISO 8601 format.
TIME	The time in ISO 8601 format.
time_t(ms)	The number of milliseconds that have passed since 1 January 1970.
NODE (CELL)	The node where the information was gathered.
NAME	The name of the share where the information was gathered.
B.READ(kB)	The amount of data read, rounded up to the nearest kilobyte.
B.WRIT(kB)	The amount of data written, rounded up to the nearest kilobyte.
ALL	The total of all file and directory operations.
F.CREA / D.CREA	The number of files/directories created.
F.OPEN	The number of files opened.
F.READ	The number of file read operations.
F.WRIT	The number of file write operations.
F.CLOS	The number of files closed.
F.DELE / D.DELE	The number of files/directories deleted.
F.ATTR / D.ATTR	The number of file/directory attribute write operations.
F.RENA / D.RENA	The number of file/directory rename or move operations.
F.SIZE	The number of file size changes.

Example 25.2. Server Statistics File (CSV format)

```
DATE,TIME,time_t(ms),CELL,NAME,B.READ(KB),B.WRIT(KB),ALL,F.CREA,F.OPEN,F.READ,F.WRIT,
F.CLOS,F.DELE,F.ATTR,F.SIZE,F.RENA,D.CREA,D.DELE,D.ATTR,D.RENA
2011-09-26,17:36:21,131...594,node-01,share-01,0,17467,1368,176,0,0,655,175,0,175,...
2011-09-26,17:41:21,131...591,node-01,share-01,0,18,24,4,0,0,5,5,0,5,5,0,0,0,0,0
2011-09-26,17:46:21,131...592,node-01,share-01,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
2011-09-26,17:51:21,131...644,node-01,share-01,0,103141,8178,1055,0,0,3886,1054,0,...
2011-09-26,17:56:21,131...592,node-01,share-01,0,147452,11743,1519,0,0,5567,1519,0...
2011-09-26,18:01:21,131...591,node-01,share-01,0,187379,14920,1929,0,0,7074,1929,0...
2011-09-26,18:06:21,131...823,node-01,share-01,0,249713,19821,2559,0,0,9415,2559,0...
2011-09-26,18:11:21,131...732,node-01,share-01,0,227314,18127,2346,0,0,8586,2346,0...
2011-09-26,18:16:21,131...593,node-01,share-01,0,218572,17455,2261,0,0,8261,2261,0...
```

(Three dots means characters or lines have been removed for presentation purposes.)

Figure 25.4. Server Statistics CSV file imported in a spreadsheet

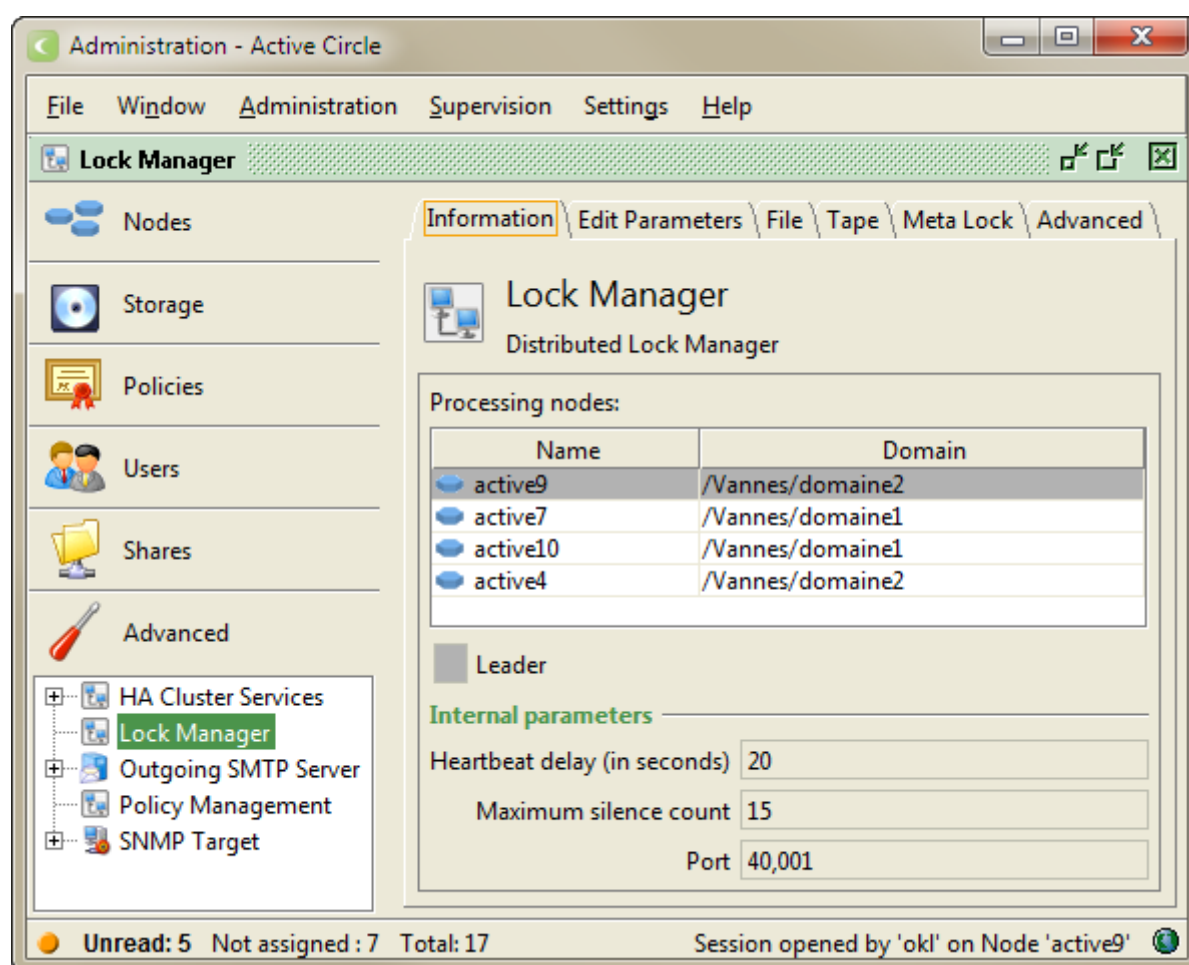
	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	DATE	TIME	time_t(ms)	CELL	NAME	B.READ(KB)	B.WRIT(KB)	ALL	F.CREA	F.OPEN	F.READ	F.WRIT	F.CLOS	F.DELE	F.ATTR	F.SIZE	F.RENA	D.CREA	D.DELE	D.ATTR	D.RENA
2	2011-09-26	17:36:21	1317051381594	node-01	share-01	0	17467	1368	176	0	0	655	175	0	175	175	0	6	0	5	1
3	2011-09-26	17:41:21	1317051681591	node-01	share-01	0	18	24	4	0	0	5	5	0	5	5	0	0	0	0	0
4	2011-09-26	17:46:21	1317051981592	node-01	share-01	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	2011-09-26	17:51:21	1317052281644	node-01	share-01	0	103141	8178	1055	0	0	3886	1054	0	1054	1055	0	37	0	36	1
6	2011-09-26	17:56:21	1317052581592	node-01	share-01	0	147452	11743	1519	0	0	5567	1519	0	1519	1519	0	50	0	50	0
7	2011-09-26	18:01:21	1317052881591	node-01	share-01	0	187379	14920	1929	0	0	7074	1929	0	1929	1929	0	65	0	65	0
8	2011-09-26	18:06:21	1317053181823	node-01	share-01	0	249713	19821	2559	0	0	9415	2559	0	2559	2559	0	85	0	85	0
9	2011-09-26	18:11:21	1317053481732	node-01	share-01	0	227314	18127	2346	0	0	8586	2346	0	2347	2346	0	78	0	78	0
10	2011-09-26	18:16:21	1317053781593	node-01	share-01	0	218572	17455	2261	0	0	8261	2261	0	2261	2261	0	75	0	75	0

Chapter 26. Distributed Lock Manager

The primary role of the Lock Manager system is to provide a system for managing concurrent access to files stored in the Active Circle system. It has been designed so that it can also be used by other internal Active Circle services that need to place locks on resources in order to guarantee integrity. This chapter describes its role in managing concurrent access.

To access this functionality in the Active Circle Administration Tool, select the **Advanced** view and click on the **Lock Manager** item.

Figure 26.1. Distributed Lock Manager Information



Important

In order for the Lock Manager to be applied to a share, it must be explicitly activated by the **Enable Distributed Lock Manager** option on the the **Edit Parameters** tab for that share. For more information, please see [“Parameters”](#) in [Chapter 23, Shares](#).

How the Lock Manager Works

From an end user's perspective, the Lock Manager functionality seems simple. When a user opens a file on a share that has the Lock Manager service activated, the system places a

lock on that file. If another user then tries to open the same file, this user will only be allowed to open the file in read-only mode. These users may be connected to shares on different nodes, as long as the lock manager is activated on each share and node.

The functionality works both for single-site environments and in multi-site environments configured with several nodes in different physical locations, as the Distributed Lock Manager functionality manages the locks across all the nodes the service has been activated on.

To illustrate, we use the following example of a simple configuration with two nodes:

Node-01 is in Paris and Node-02 is in London

There is one share, `share-01`, with the Lock Manager service activated

The users in each city connect to the share through their local node

This configuration allows users at different sites access to the same files while also ensuring the integrity of their work by guaranteeing that only one user at a time can work on any given file.



Important

Because of the way in which certain applications manage how they open files, the file locking mechanism does not work with all file formats. In order for the Lock Manager to function, the application used to modify the file must maintain a connection to the file itself during modification. This is the case for common office productivity software, such as Open Office and Microsoft Office. It does not, however, apply to applications like Microsoft's Notepad.

Lock Manager Administration

When you select **Lock Manager** in the **Advanced** view, several tabs are available which allow you to perform configuration and administration tasks:

- [Processing nodes and internal parameters](#)
- [List of locks](#)
- [Advanced information](#)

Nodes and Parameters

The **Edit Parameters** tab allows you to configure the nodes and internal parameters for the Lock Manager service.

Figure 26.2. Edit Parameters tab

Information **Edit Parameters** File Tape Meta Lock Advanced

Last update: May 23, 2013 11:43:04 AM

Processing nodes:

Add Remove Refresh

Name	Domain	Network Interfaces
active9	/Vannes/domaine2	<AUTO DETECTION>
active7	/Vannes/domaine1	<AUTO DETECTION>
active10	/Vannes/domaine1	<AUTO DETECTION>
active4	/Vannes/domaine2	<AUTO DETECTION>

Internal parameters

Heartbeat delay (in seconds) 2

Maximum silence count 15

Port 40,001

Apply Cancel

The list of **Processing nodes** shows which nodes are part of the lock manager cluster, meaning that they are eligible for being elected *leader*. The leader node takes on the *Dispatcher* role, and it communicates this to all the nodes in the Circle. The leader node is indicated by a grey background in the list of nodes. The other nodes will then pass all lock requests to this node. The dispatcher node keeps a directory of all the resource lock managers in the Circle, enabling it to manage the lock requests of the Circle. The next node on this list is the secondary dispatcher, which keeps copies of the localization information for the lock managers distributed by the lead dispatcher. In a failover situation, if the the lead node becomes unavailable, the secondary dispatcher will take over the lead role, having already the necessary data for maintaining the lock manager directory.

Use the up and down arrows to the right of the **Processing nodes** list to change the preferred order used in the election of the lead dispatcher node. Use the **Add** and **Remove** buttons above the list to add or remove nodes from the list.

Click on the **Refresh** button to update the display of the processing node list.



Note

The lock manager functionality is not limited to the nodes on the list of **Processing nodes**. The functionality applies to all shares for which the Distributed Lock Manager is enabled. Consequently, any node where such a share is mounted will be included. A Resource Lock Manager (RLM) on an object can be hosted by any node, but the lead dispatcher node handles the requests and distributes the RLMs. For each RLM residing on a node, there can be a secondary RLM on another node which is a candidate to take over the RLM function for the object in case of failover of the node hosting the RLM.

Use the settings under **Internal parameters** to configure settings for the communication between the nodes.

Heartbeat delay (in seconds)	The frequency of messages sent from the node.
Maximum silence count	Maximum number of lost UDP messages allowed from a node before the other nodes in the cluster consider that it is down and that its workload must be shifted to another node.
Port	The port number on which the nodes are listening for heartbeats (40001 by default).

Viewing Lock Lists

There are several types of locks managed by the Lock Manager, each listed under a separate tab:

- **File:** File locks, including consolidated locks
- **Tape:** Locks on tapes, drives and tape libraries
- **Cloud:** Locks on archived resources in cloud pools
- **Meta Lock:** Other locks in the Circle, such as on the catalogs

The lock lists on the tabs display the current active locks in the Circle and information concerning them, such as the name of the locked file, resource type, lock type, the name of the user who owns the lock, and the node currently hosting the Lock Manager. You can add or remove information columns by right-clicking on a column heading and selecting **Column properties....** A dialog will be displayed where you can select and deselect columns for display.

Figure 26.3. Lock List tabs

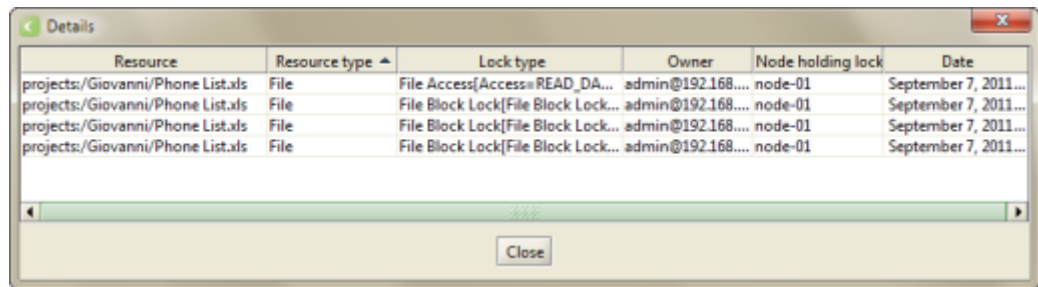
Resource	Resource type	Lock type	Owner	Node holding lock	Date
partage/Finance	File	File Access(Access=READ_DATA, Sharing...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:56:3
partage/Finance/Factures 2013	File	File Access(Access=READ_DATA, Sharing...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:56:4
partage/Finance/Factures 2013/Fichier ...	File	File Access(Access=READ_DATA, Sharing...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:56:4
partage/Finance/Factures 2013/Fichier ...	File	File Access(Access=READ_DATA + WRITE...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:56:4
partage/Finance/Factures 2013/Fichier ...	File	File Access(Access=READ_DATA + WRITE...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:56:4
partage/Ressources Humaines	File	File Access(Access=READ_DATA, Sharing...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:12:3
partage/Ressources Humaines/Congés...	File	File Access(Access=READ_DATA + WRITE...	xli@192.168.2.136	jouycircle1	May 23, 2013 10:12:4
partage/Ressources Humaines/Congés...	File	File Access(Access=READ_DATA + WRITE...	xli@192.168.2.136	jouycircle1	May 23, 2013 11:47:3
partage/Utilisateurs/okl	File	File Access(Access=READ_DATA, Sharing...	okleven@192.16...	vannescircle1	May 23, 2013 12:09:3
partage/Utilisateurs/okl/working	File	File Access(Access=READ_DATA, Sharing...	okleven@192.16...	vannescircle1	May 23, 2013 12:09:3
partage/Utilisateurs/okl/working/Doc ...	File	Consolidated	okleven@192.16...	vannescircle1	May 23, 2013 12:09:3

The **Release** button provides you with a mechanism to unlock a file that has been locked by a user. This might be used in the event another user urgently needs access to the file.

The **Refresh** button allows you to update the lock information displayed in the list.

Some locks are composed of several individual locks (for example files opened by Microsoft Word). Such locks are grouped into one line with **Consolidated** indicated in the **Lock type** column. For a consolidated lock you can select **Details** to display a window with more information on the lock.

Figure 26.4. Lock Details

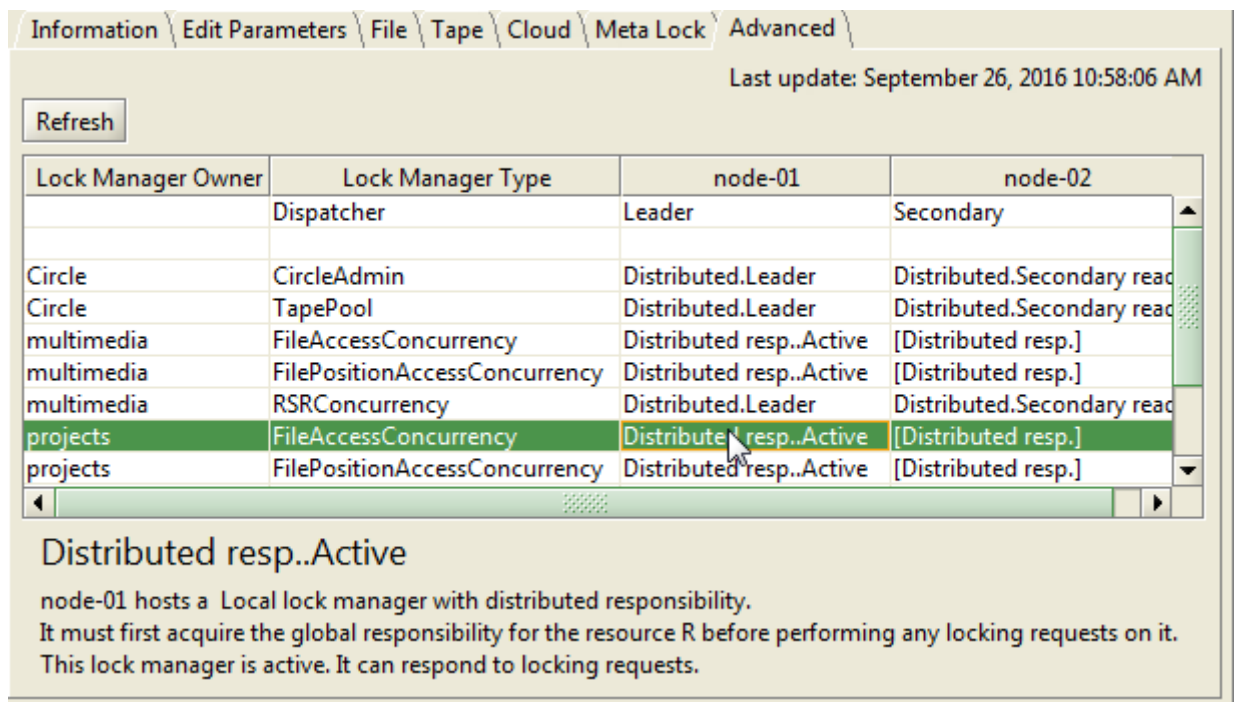


Resource	Resource type	Lock type	Owner	Node holding lock	Date
projects:/Giovanni/Phone List.xls	File	File Access[Access=READ_DA...	admin@192.168...	node-01	September 7, 2011...
projects:/Giovanni/Phone List.xls	File	File Block Lock[File Block Lock...	admin@192.168...	node-01	September 7, 2011...
projects:/Giovanni/Phone List.xls	File	File Block Lock[File Block Lock...	admin@192.168...	node-01	September 7, 2011...
projects:/Giovanni/Phone List.xls	File	File Block Lock[File Block Lock...	admin@192.168...	node-01	September 7, 2011...

Advanced

The **Advanced** tab provides information about the current state of the Lock Manager components. Each cell in the table contains more information which will be displayed at the bottom of the dialog box when you hover the mouse pointer over (or click in) the cell. This gives you detailed information on the Lock Managers in the Circle. It shows which node is the lead dispatcher, which nodes host secondary lock managers, which nodes are candidates for hosting a lock manager, and so on. *Candidate* status is indicated by [square brackets] surrounding the cell entry.

Figure 26.5. Advanced tab – Local Lock Manager



Lock Manager Owner	Lock Manager Type	node-01	node-02
	Dispatcher	Leader	Secondary
Circle	CircleAdmin	Distributed.Leader	Distributed.Secondary read
Circle	TapePool	Distributed.Leader	Distributed.Secondary read
multimedia	FileAccessConcurrency	Distributed resp..Active	[Distributed resp.]
multimedia	FilePositionAccessConcurrency	Distributed resp..Active	[Distributed resp.]
multimedia	RSRConcurrency	Distributed.Leader	Distributed.Secondary read
projects	FileAccessConcurrency	Distributed resp..Active	[Distributed resp.]
projects	FilePositionAccessConcurrency	Distributed resp..Active	[Distributed resp.]

Distributed resp..Active
 node-01 hosts a Local lock manager with distributed responsibility.
 It must first acquire the global responsibility for the resource R before performing any locking requests on it.
 This lock manager is active. It can respond to locking requests.

Figure 26.6. Advanced tab – Distributed Lock Manager

Information | Edit Parameters | File | Tape | Cloud | Meta Lock | **Advanced**

Last update: September 26, 2016 10:58:06 AM

Refresh

Lock Manager Owner	Lock Manager Type	node-01	node-02
	Dispatcher	Leader	Secondary
	CircleAdmin	Distributed.Leader	Distributed.Secondary ready
	TapePool	Distributed.Leader	Distributed.Secondary ready
media	FileAccessConcurrency	Distributed resp..Active	[Distributed resp.]
media	FilePositionAccessConcurrency	Distributed resp..Active	[Distributed resp.]
media	RSRConcurrency	Distributed.Leader	Distributed.Secondary ready
cts	FileAccessConcurrency	Distributed resp..Active	[Distributed resp.]
cts	FilePositionAccessConcurrency	Distributed resp..Active	[Distributed resp.]

Distributed.Secondary ready

node-02 hosts a Distributed lock manager.
Only one node can lock the resource R in the Circle.
This lock manager only keeps a single copy of locks verified by the leader.

Conflict Resolution

There could be situations where locks are unresolved or inconsistent with the information in the lock manager administration interface as described above. Reasons could be network failures, server failures, lock manager cluster leader failures or other technical problems. In normal circumstances the Active Circle system is able to handle this, but unforeseen situations might arise where resources become inaccessible due to unresolved or conflicting locks being held at the nodes. To avoid having to stop entire nodes to release or resolve locks, it is possible to stop (and restart) the distributed lock manager (DRLM) services on the node in question. This is done using the **Controller of DRLM services** component for the node on the “Advanced” tab in the Administration Tool **Supervision** interface (see [Chapter 32, Supervision](#)).



Important

Stopping the DRLM controller services is a temporary measure to allow the other nodes in the circle to recover control of locks. The services should be restarted when the problems are resolved. Care should be taken if attempting to manipulate lock manager operation, as it might result in unforeseen consequences. It is recommended to contact Active Circle product support for assistance with dealing with such scenarios.

Chapter 27. Policy Management

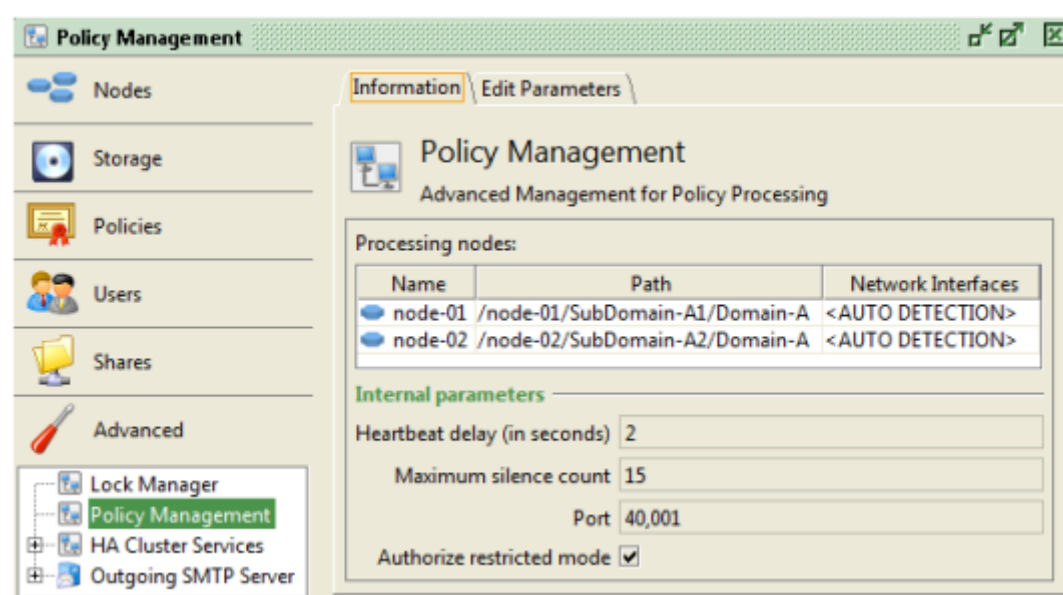
Using the Policy Management service, it is possible to create a "policy cluster", that is a list of nodes that together are responsible for processing the policies. The first node in the list distributes the processing to the other nodes in the list. If the first node is not available, the second node in the list takes over, and so on. Using Active Circle Policy Management, you can implement a robust strategy for guaranteeing continuity of service.

Using Policy Management

In order for a node to be able to participate in a policy processing cluster, the Policy Management service must be enabled for that node. This option is activated by the administrator during the creation of the node (see ["Adding a New Node"](#)). If policy management is enabled for a node it will appear in the list of nodes.

To view the nodes with policy management enabled, click on **Policy Management** in the **Policies** view:

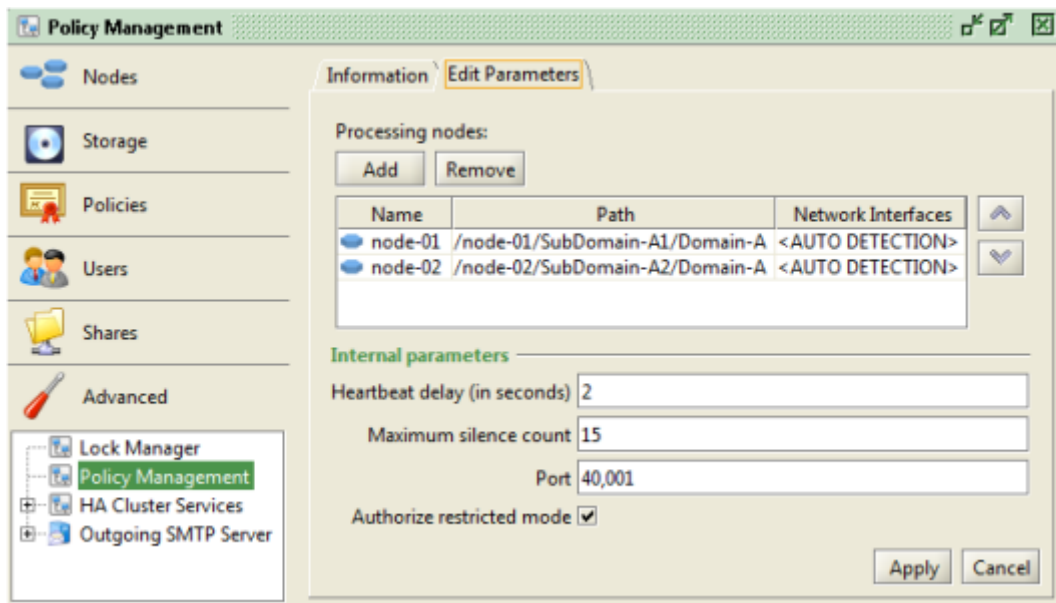
Figure 27.1. Policy Management Information tab



Edit Parameters

You can configure the Policy Management cluster in the **Edit Parameters** tab.

Figure 27.2. Edit Parameters tab



By selecting a node in the list and clicking the **Up** or **Down** arrow buttons, you can change the order of the nodes in the cluster.

Once a node in the list is assigned to manage the policies for a share, it will remain associated with that share until the Active Circle service on that node is restarted.

If one of the nodes becomes unavailable, the first node chooses one from the remaining nodes in the list to guarantee the services impacted by the change. The policy management algorithm has been designed to distribute the policy processing for the various shares across all the nodes in the cluster.

You can configure the following **Internal parameters**:

Heartbeat delay (in seconds)	The frequency of messages sent (in Multicast) from each of the nodes in the Policy Management cluster to the other nodes to indicate their availability for policy processing.
Maximum silence count	Maximum number of lost UDP messages allowed from a node before the other nodes in the cluster consider that it is down and that its workload must be shifted to another policy processing node.
Port	The port number on which the nodes are listening for heartbeats (40001 by default).
Authorize restricted mode	This option allows the system to reduce the load on a node if it is the only one left in the cluster. If this occurs and restricted mode is authorized (it is by default), then only new data created on a share will be processed according to their policies. Existing data will not be analyzed.

Chapter 28. HA Cluster Services

The **HA Cluster Services** (or High Availability Cluster Services) item in the **Advanced View** allows administrators to create a virtual file server consisting of two or more nodes in order to provide users with a fault-free point of connection to the Virtual File System. For example, a user can map a drive to a virtual file server with a fixed IP address managed by the Active Circle system, instead of to an IP address or name associated with a specific machine.

Consequently, the system administrator can stop one of the nodes in an HA Cluster (for example to perform maintenance on the server) without interrupting the service.

The architecture of the Active Circle system has been designed to ensure the security of your data. With the creation of a HA Cluster, it is possible to provide a continuity of service that virtually guarantees users uninterrupted access to their data.

A High Availability Cluster is based on the notion of a group of nodes that together provide the redundancy required to guarantee continuity of service. The cluster functions like a virtual file server to which users can map a network drive. If the node in the cluster that provides users access to the system suddenly stops working, the other nodes in the group detect the problem and one of them is chosen to replace it.

Because the users are mapped to a virtual file server and not a real one, they do not need to manually switch over to another machine in order to continue working.

The implementation of an HA Cluster requires knowledge of how the network in which it is to be implemented is configured, in particular how IP addresses are configured, available IP ranges, and the network topology.

Additionally, because the solution is based on IP, UDP, Ethernet, and ARP technologies, some familiarity with these is also required.

Election

When an HA Cluster is started, the nodes that make up the cluster elect one node to respond to any requests that are sent to the cluster's IP address. The method used for choosing which node in the HA Cluster will provide users access to the shares associated with the cluster is based on an election process. The nodes communicate amongst each other through multicast UDP, with the lead node signalling its presence to the others through its "heartbeat" (by default once every three seconds).

If the lead node in the HA Cluster stops emitting its heartbeat, the other nodes detect this. If the heartbeat is still not detected after a certain number of missed signals, then the other nodes initiate an election.

Each of the nodes in the cluster broadcast their candidacy to the others. When a node detects message from a node with a higher priority (determined by its position in the list of the cluster's nodes), then that node considers that it has lost the election. If a node does not receive a signal from any other node with a higher priority within a certain time limit, then the node considers it has won the election and takes over the IP address.

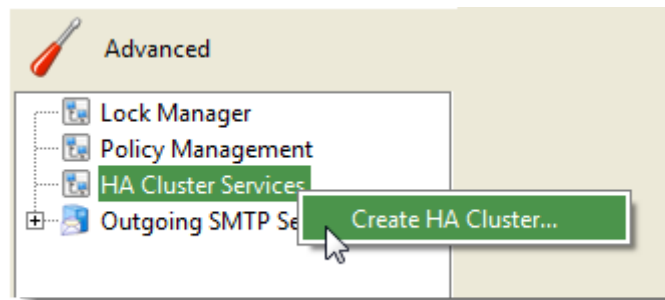
The sections below describe how to create HA Clusters and the different methods for configuring them.

Creating and Configuring HA Clusters

There are four main tasks involved in creating and configuring an HA Cluster:

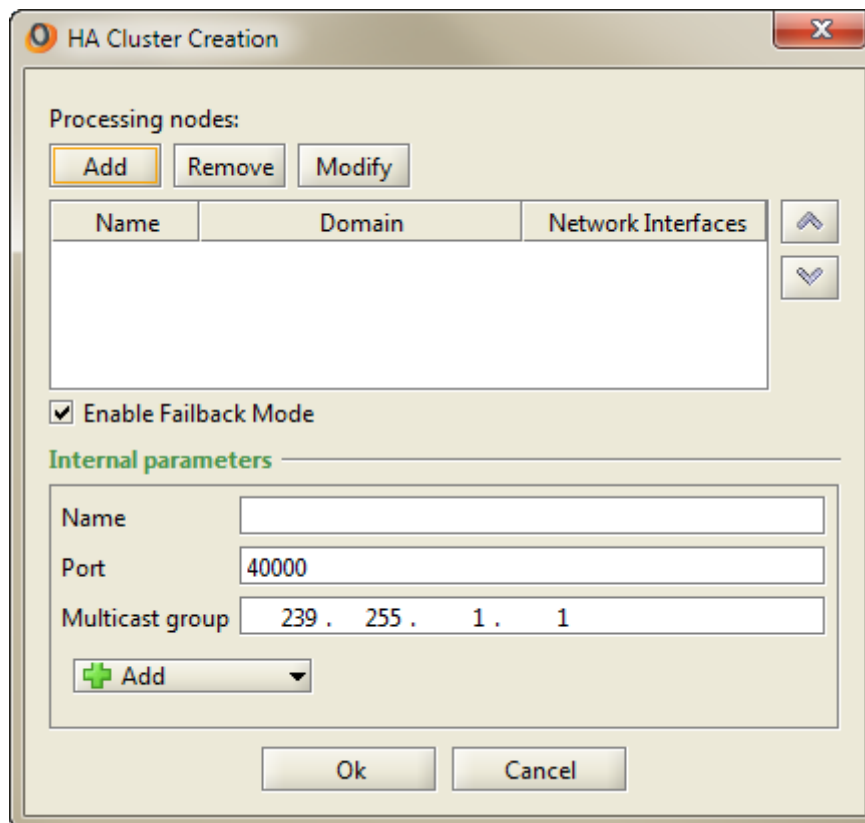
- [Adding processing nodes for the cluster](#)
- [Defining node priority](#)
- [Specifying internal network parameters](#)
- [Selecting and configuring the cluster type](#)

To create an HA Cluster, right-click on the **HA Cluster Services** item in the **Advanced** view and select **Create HA Cluster**:



The **HA Cluster Creation** dialog box appears:


Figure 28.1. HA Cluster Creation Dialog



Adding Processing Nodes

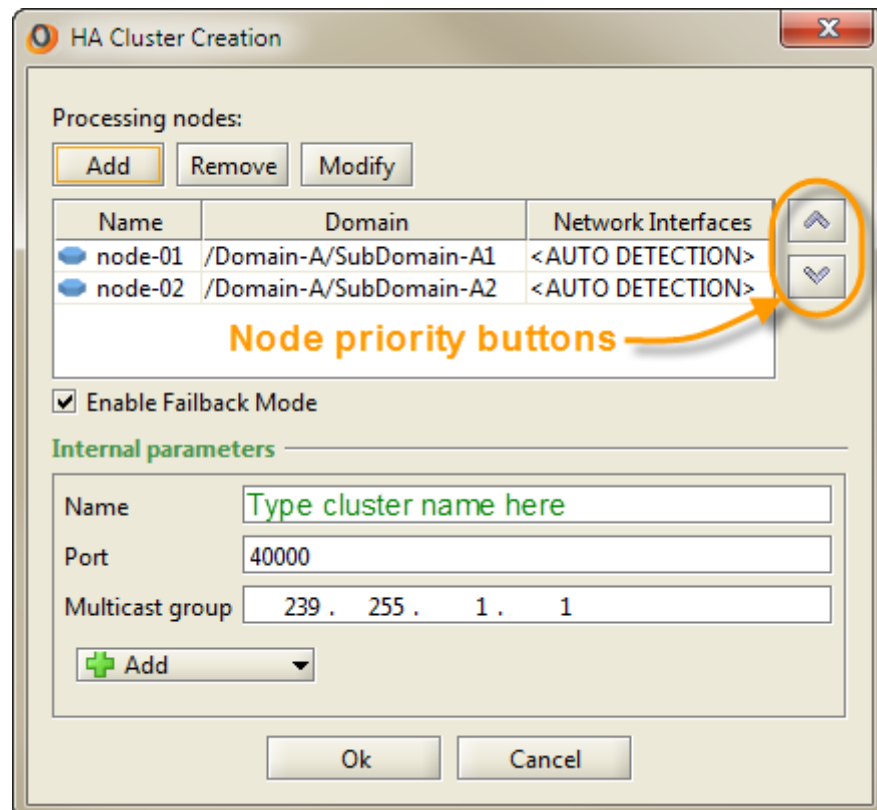
The first step is to add the nodes to your cluster. You must associate at least two nodes with it. Click the **Add** button under the heading **Processing nodes** to display the **Add Nodes** dialog box:

Figure 28.2. Add Nodes



Select the nodes you wish to assign to the cluster. Any nodes already assigned to the current cluster will appear greyed out. Click the **OK** button to add the nodes to the processing nodes table.

Figure 28.3. Nodes Added



Node Priority

You can modify the priority of the nodes in the list by using the arrow buttons on the right-hand side of the table. The first node in the processing node list (the preferred 'leader node') is the one that will be assigned the role of providing users access to the system. If that node fails or is taken offline for maintenance, an election is initiated and the second node (if it is available) will then be assigned the role because it has the highest priority. When the node at the top of the list comes back in service, it will cause a re-election to take place so that it can again become the leader. The leader node line is indicated by gray shading in the list. The priority of the nodes can be disabled using the [failback mode](#) option. If the failback mode has been disabled, the arrow buttons will be deactivated (greyed out).

The node priority settings can be modified after the HA cluster is created, by using the **Edit Parameters** tab for the cluster.

Failback Mode

By default, a cluster will always attempt to fail back to the leader node as defined by the [node priority](#). You can disable this functionality by removing the check mark in the **Enable Failback Mode** box. When the automatic failback mode is disabled, the order of the nodes in the **Processing nodes** will still be respected in the case of a re-election following the failure of a leader node. However, if a previous leader node or the node at the top of the list comes back online after downtime, a re-election will not be triggered as long as the current leader node is in service, unless you change the order of the nodes or click on the [Restart](#) button for the cluster.

When automatic failback is disabled, the system will issue a supervision note as information each time a new leader node is elected.

The node failback mode can be modified after the HA cluster is created, by using the **Edit Parameters** tab for the cluster.

General Settings

Specify the general settings for the cluster under **Internal parameters**:

Name	Provide a name to identify the HA cluster. Note that this name is not the name used to identify the virtual file server over the network (but the name can be the same).
Port	The UDP port on which the nodes are listening for messages from other members of the cluster. <i>Default value: 40000</i>
Multicast group	Destination IP address for messages sent by the nodes. <i>Default value: 239.255.1.1</i>

Cluster Types

When you have selected the nodes to include in the cluster and the general parameters, you must select and configure the type of cluster you would like to create. There are four types of HA Clusters:

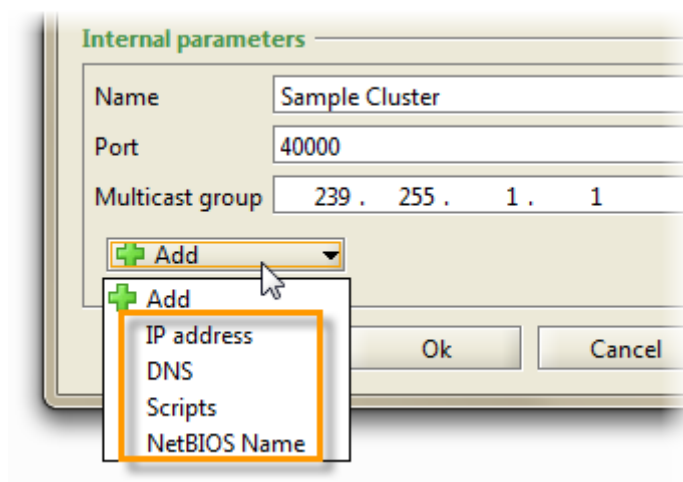
IP Address	The simplest method for creating an HA Cluster is defining an IP address for the virtual file server to which users will map their
-------------------	--

connections. Due to routing considerations, the nodes of an IP cluster must be in the same LAN.

- DNS Name** A DNS cluster allows you to declare a host name for the cluster. This is only available for systems running BIND9.
- Scripts** This option allows you to configure your cluster to call scripts that create a VPN tunnel that can be used as access point.
- NetBIOS Name** This option allows you to declare a NetBIOS Name for the cluster. Users can then use this name instead of, or in addition to, an IP address or a hostname (for example if DNS is not enabled for the network).

You specify the cluster access method(s) to use with the cluster by clicking on the **+ Add** button under **Internal parameters**:

Figure 28.4. Add Cluster Type



You must decide which method (or methods) to use to connect to your cluster. The procedures for configuring each type of cluster are described in the sections below. You will not be able to create the cluster before you have completed the procedure for at least one of the methods (you will receive a message reminding you of this if you click on the **OK** button at this point).



Note

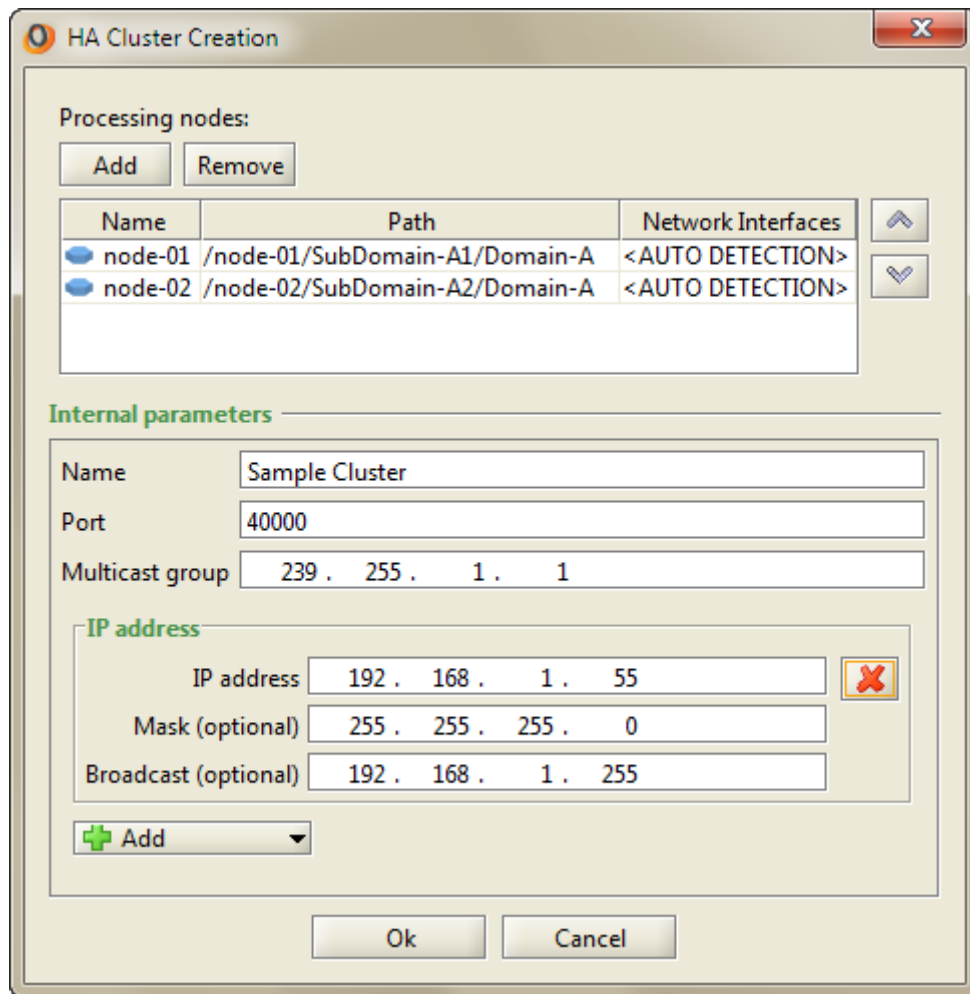
The different types of HA Clusters are not mutually exclusive. For example, you can create a cluster and assign it both an IP address and a domain name.

IP Cluster

Click on the **+ Add** button under **Internal parameters** and select **IP address**.

Selecting the **IP Address** option will add fields to the dialog box to allow you to configure the IP Cluster (to see the added fields, you may need to use the scroll bar or drag the window border to make the dialog larger).

Figure 28.5. Configure an IP Cluster



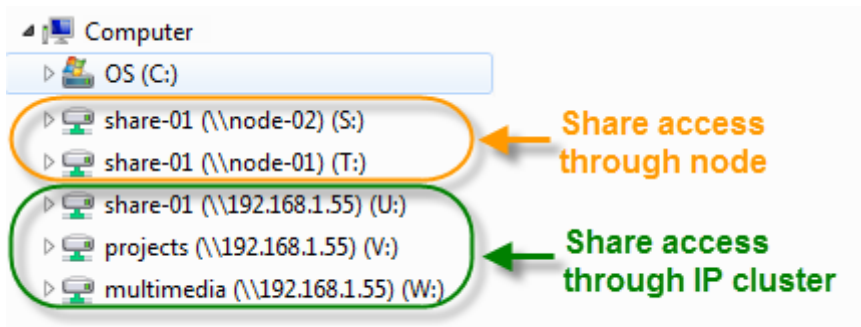
IP Address The IP address for the virtual file server. It is to this address that users will map their network connections (instead of to the actual IP address of one of the Circle servers) in order to access the files stored on the system.

Mask The subnet mask (optional)

Broadcast The broadcast mask (optional)

Click **OK** to create the cluster and close the dialog box. When created, clients can access the Circle using the cluster IP address instead of mounting shares on individual nodes.

Figure 28.6. Windows Client Access to Shares in HA Cluster



You can modify the HA Cluster at any time by selecting it in the **Advanced** view and clicking on the **Edit Parameters** tab (see “[Edit Parameters](#)”).

DNS Cluster

In addition to allowing you to declare an IP address for your virtual file server, you can also provide it with a domain name by configuring a DNS cluster. Active Circle supports BIND9 as the DNS server to use for configuring a DNS cluster.



Note

The different types of HA Clusters are not mutually exclusive. For example, you can create a cluster and assign it both an IP address and a domain name.

Because the creation of a DNS cluster involves making a change to the configuration of the DNS server, all communication between the DNS client (in this case, an Active Circle node) and the DNS server requires a key to validate the communications between client and server. Additional information on how to configure the DNS server is provided in [Appendix C, Configuring external software for use with HA Clusters](#).

To modify an existing HA Cluster, select it in the **Advanced** view and click on the **Edit Parameters** tab. To configure DNS for an HA Cluster, click on the **+ Add** button under **Internal parameters** and select **DNS**. This will add a **DNS** section to the **Edit Parameters** tab (or to the **HA Cluster Creation** dialog box when creating a new HA cluster).

Figure 28.7. Create a DNS Cluster

Host name to declare

The name of the virtual file server to be declared on the DNS server.

Domain

Name of the domain to which the host will be added.

For example: *active-circle.com*.

TSIG key name	A unique name used to identify the TSIG (Transaction SIGnature) key. This is the name used to identify the key in the BIND9 DNS configuration.
TSIG Key (HMAC-MD5 base64)	The alphanumeric key generated with the <code>dnssec-keygen</code> command. For more information on how to configure BIND9, please refer to Appendix C, Configuring external software for use with HA Clusters .
DNS Server list	The list of the IP addresses of the DNS servers on which to register the new DNS name.
Addresses to favor	A list of preferred IP addresses to use if any of the nodes in the clusters have multiple network adapters. This feature allows you to manage which network adapters are used for the different types of network traffic. When a node is elected leader of a cluster, it looks at its available IP addresses (ignoring the loopback address). By default, it will use the first address it finds. However, if for a given node a preferred address is defined here, then that address (and the corresponding network card) will be used.

Click on the **Apply** button (or the **OK** button if you are creating a new cluster) to save the configuration changes.

Scripted Cluster with OpenVPN

A third option for creating an HA Cluster is by using scripts to launch a VPN service on the lead node of the cluster that provides the virtual network interface to which users will connect.

Active Circle supports the use of [OpenVPN](#), an open source VPN solution that is available for GNU/Linux, Sun Solaris and Microsoft Windows. In order to implement a scripted cluster, OpenVPN must be installed on each node in the HA cluster as well as on any client machines that need to access the cluster via this connection. For more information on configuring OpenVPN, see [“Configuring Scripted Clusters with OpenVPN”](#).

As with the other methods for configuring HA clusters, you can configure a cluster with scripts when you create the cluster or by going to the **Edit Parameters** tab for an existing cluster and clicking on the **+ Add** button and selecting **Scripts**.

Figure 28.8. Create an OpenVPN Cluster

Scripts		
Script for taking control of the resource	<input type="text" value="startOpenVPN"/>	parameters <input type="text"/>
Script for confirming control of the resource	<input type="text" value="restartOpenVPN"/>	parameters <input type="text"/>
Script for giving up control of the resource	<input type="text" value="stopOpenVPN"/>	parameters <input type="text"/>

As in the case of the other cluster types, you must provide a name, a port number, and an IP address for the Multicast group. All clusters must have at least two nodes associated with it.

There are three lines where you must provide the names of the scripts used to manage OpenVPN.

- Script for taking control of the resource. Used to start the OpenVPN service.
- Script for confirming control of the resource. Used to stop and then re-start the OpenVPN service.
- Script for giving up control of the resource. Used to stop the OpenVPN service.

To complete the creation of the cluster, you must associate at least two nodes with it. Refer to the procedure described in the [Add Nodes](#) section above.

NetBIOS Cluster

To accommodate clients using NetBIOS, you can configure a NetBIOS Name as the identifier for the HA Cluster. The name may be configured using a WINS Server in the network, or by using simple broadcast between the clients and nodes. Note that if the NetBIOS name is not configured for the whole network through a WINS Server, the NetBIOS name discovery will only be possible for clients that belong to the same subnet as the Active Circle nodes. If no WINS Server is used, no pre-configuration is necessary. Just create a unique name of up to 15 characters, and the Active Circle nodes on the subnet will automatically use it for broadcast and cluster election. You can add a NetBIOS name to the HA Cluster in addition to an IP address or a DNS configuration.


You can add configuration for a NetBIOS cluster when you create the cluster or by going to the **Edit Parameters** tab for an existing cluster and clicking on the **+ Add** button and selecting **NetBIOS**.

Figure 28.9. Create a NetBIOS Cluster



Click on **Apply** (or **OK** if you are creating a new cluster) to add or create the NetBIOS configuration for your HA Cluster. Wait for the processing to finish. After creation, you can mount the shares that have been associated with the cluster using the NetBIOS name.

Figure 28.10. Accessing Clusters using Windows Drive Mapping



Edit Parameters

In the **Edit Parameters** tab for an HA Cluster you can modify any of the settings for a selected cluster. The tab is organized into a **Processing nodes** area at the top, followed by the

Internal Parameters area that contains separate sections for the access modes which have been configured for the cluster.

Using this tab, you can perform the following tasks:

- **Add** nodes to the list of processing nodes
- **Remove** nodes from the list
- **Modify** the network interfaces
- **Restart** the HA cluster
- **Refresh** the leader node information on the tab
- Modify the node priority and cluster failback settings
- Edit the cluster name and description
- Modify network parameters (heartbeat, silence count, port and multicast group)
- Add or modify cluster access modes (IP address, NetBIOS name, etc.)

Figure 28.11. Edit Parameters tab

Information **Edit Parameters** Shares

Last update: June 1, 2012 11:41:05 AM

Processing nodes:

Add Remove Modify Restart Refresh

Name	Domain	Network Interfaces
node-02	/Domain-A/SubDomain-A2	<AUTO DETECTION>
node-01	/Domain-A/SubDomain-A1	<AUTO DETECTION>
vcentos	/Domain-A/SubDomain-A1	<AUTO DETECTION>

☒ Enable Failback Mode

Internal parameters

Name: Doc Cluster

Description: Test cluster for Doc dept.

Heartbeat delay (in seconds): 3

Maximum silence count: 15

Port: 40,000

Multicast group: 239 . 255 . 1 . 1

IP address

IP address: 192 . 168 . 1 . 90

Mask (optional): 255 . 255 . 255 . 0

Broadcast (optional): 192 . 168 . 1 . 255

NetBIOS Name

NetBIOS Name: DOC-HACLUSTER

+ Add

Apply Cancel

Node Priority and Failback Mode

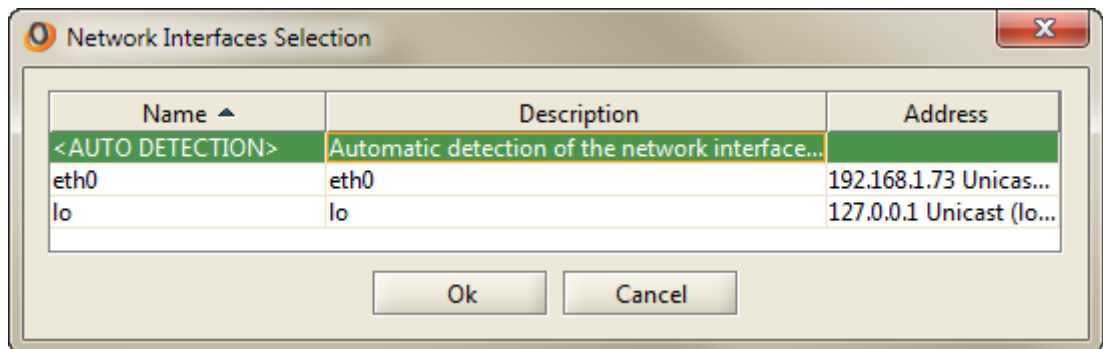
Using the arrow buttons to the right of the **Processing nodes** list and the **Enable Failback Mode** check box below the list, you can modify the setting for these features. For an explanation, please see [this topic](#) under “Creating and Configuring HA Clusters”.

Network Interface

If the nodes in the cluster possess only a single network adapter, then the label **<AUTOMATIC DETECTION>** will appear in the list of nodes. If a node possesses more than one network adapter, then the first one detected (ignoring the loopback address) will be selected.

You can change which network adapter is used by a node by selecting it in the list and clicking the **Modify** button and selecting the appropriate adapter in the list.

Figure 28.12. Network Interface Selection



Network Settings

Heartbeat delay (in seconds)	The frequency with which the lead node in the cluster transmits a signal to the other members to indicate that it is functioning properly.
Maximum silence count	The number of missed "heartbeats" after which the other nodes in the cluster initiate an election to select a new lead node.
Port	UDP port on which the nodes in the cluster listen for "heartbeats" from the lead node.
Multicast group	The IP address to which the "heartbeats" are sent.

Share dependency

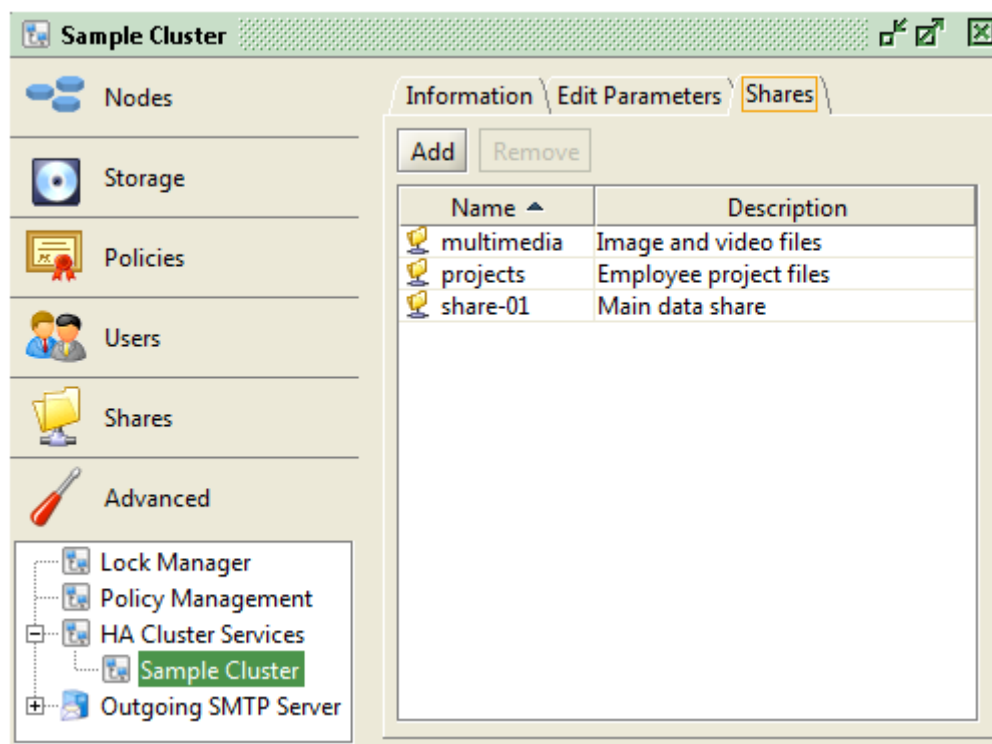
When an HA Cluster is configured and started, all the shares associated with the nodes in the cluster will automatically be accessible through the cluster IP address or name. It is not necessary to specifically add shares to the cluster for user access purposes. However, to ensure that access to a share is maintained during leader node transfer and share synchronization, the share needs to be added to the list in the **Shares** tab for the selected cluster. Adding a share dependency here will prevent the share access from being transferred to a new leader node before that node is ready to make the share available.

One aspect of this functionality can be illustrated by the following example. Imagine a circle consisting of node A and node B configured as an IP cluster in normal operation. Node A is at

the top of the processing nodes list. It is therefore the leader, and the IP address of the cluster is held by this node. At a point in time, node A goes down. The cluster automatically goes through the election process, and node B is given the IP address and will act as leader to host the share access. While node A is down, the users continue to work on the shares through node B. After a while node A comes back online, which triggers another election process. Node A will be elected leader, but before this node can give access to the shares, it must synchronize its (now backlevel) data with the share data which have been modified during the downtime. Such synchronization could take an hour or more with large data volumes or an extended downtime. If the cluster IP and shares are allowed to be hosted by node A before this synchronization is completed, the shares will be inaccessible during the whole synchronization process. Adding the share(s) to the list of associated shares will prevent this transfer from happening immediately. In this case, access to the shares on the list will not be transferred from node B until node A has finished synchronizing and is ready to provide uninterrupted user access to the shares.

To add a share dependency for an HA cluster, click the **Add** button in the **Shares** tab and select the appropriate share(s) in the list. Click **OK** to close the dialog box.

Figure 28.13. HA Cluster – Shares tab



Chapter 29. Outgoing SMTP Server

Active Circle Nodes can be configured to send e-mail notifications to selected users in response to system events. You configure one or more e-mail accounts using the **Outgoing SMTP Server** item.

Ideally, a dedicated e-mail account should be created for this purpose.

You can either configure individual accounts for each node or domain, or configure a single account for the entire system.

Configure an E-mail Account

To configure an e-mail account, select either the **Outgoing SMTP Server** item or an individual node under this item in the **Advanced** view. Go to the **Edit Parameters** tab and enter the information for the account.

Figure 29.1. Edit Parameters tab

The screenshot shows a configuration window titled '*node-01'. On the left is a sidebar with icons for Nodes, Storage, Policies, Users, Shares, and Advanced. Under 'Advanced', there is a tree view showing 'Lock Manager', 'Policy Management', 'HA Cluster Services', and 'Outgoing SMTP Server'. Under 'Outgoing SMTP Server', 'node-01' and 'node-02' are listed, with 'node-01' selected. The main area has two tabs: 'Information' and 'Edit Parameters'. The 'Edit Parameters' tab is active and contains the following fields:

- Server address: ☐
- Port: ☐
- Sender e-mail address: ☐
- Use authentication: ☒
- User name: ☐

At the bottom right are three buttons: , , and .

Server address

Address of the outgoing e-mail server. For example, *smtp.active-circle.com*

Port

E-mail server port number. The standard is port number 25.

Sender e-mail address

The address of the sender, for example the cell.

Use authentication

Check box. This must be activated if the server requires authentication for sending e-mails. If activated, you must provide a password using the **Modify password** button.

User name

The user name associated with the e-mail address.

There are check boxes following the fields for **Server address**, **Port**, **User name**, and **Sender e-mail address**. If the root element — in this case the **Outgoing SMTP Server** item — has been configured, you can check these boxes to cause the nodes to inherit their values for these fields from the root.

Chapter 29. Outgoing SMTP Server

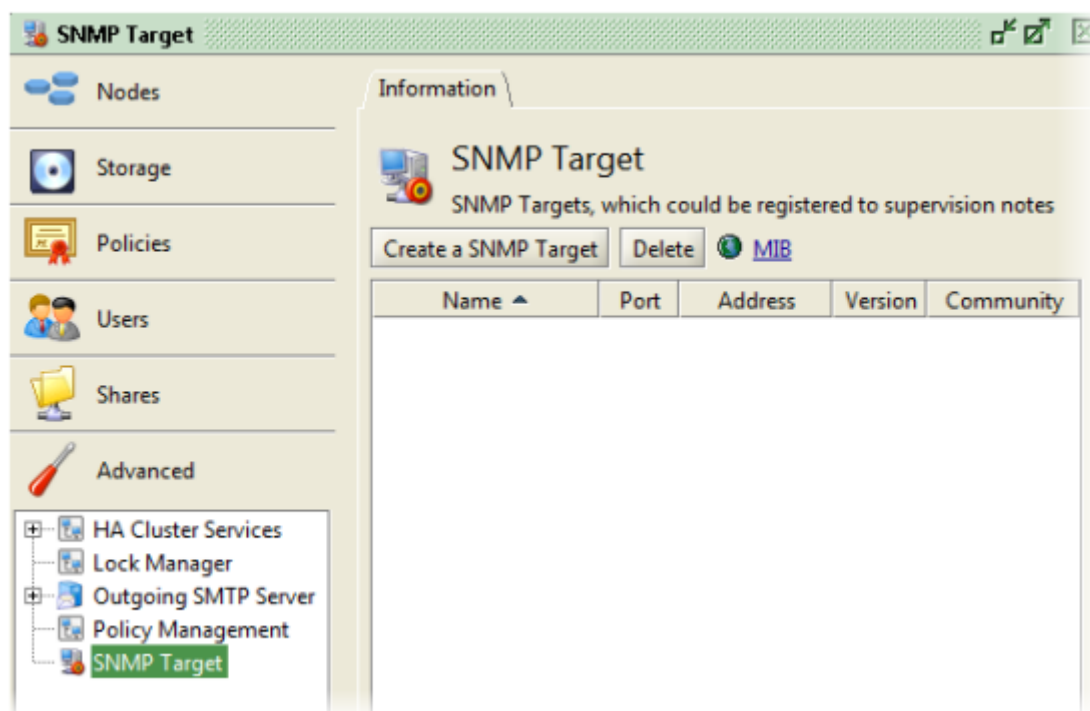
You can verify the configuration of the account by using the **Test** button. Enter an e-mail address (for example your own) and click **OK** to send the e-mail.

Chapter 30. SNMP Target

Active Circle lets you to configure a target for supervision notes to be sent to third-party systems using the Simple Network Management Protocol. You can also download the MIB file (Management Information Base) for SNMP trap definitions.

To access this functionality, go to the **Advanced** view and select the **SNMP Target** item.

Figure 30.1. SNMP Target



The packets sent via *SNMP* contains several information items, each represented by a unique object identifier (OID). The base OID is 1.3.6.1.4.1.21399.1.1.2.1. (which is registered at IANA and uniquely identifies the Active Circle application), followed by the identifier for the field:

Table 30.1. SNMP Data

SNMP Object ID	Field
1	Name
2	Date
3	Source
4	Level
5	Description
6	Cause

There is currently no mechanism in place for receiving *SNMP* messages from external sources.

Create SNMP Target

To create an **SNMP target**, click the **Create** button. The **Create SNMP Target** dialog box appears.

Figure 30.2. Create SNMP target dialog box

Name	A name to identify the target.
Description	A description to help identify the target.
IP Address	The IP address of the remote system.
Port	The port for the SNMP service on the remote system. The default SNMP port is 162.
Protocol version	The version of the SNMP protocol to use. Currently supported versions are 1 and 2c (the default).
Community	This allows the client (the Active Circle system in this instance) to identify itself to the remote system, functioning something like a password.

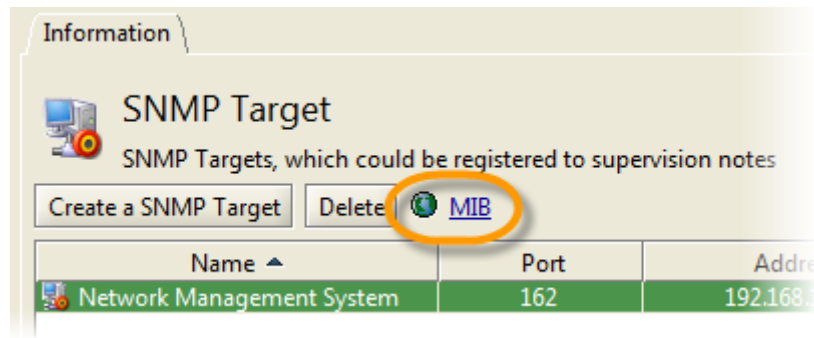
Click on the **Create** button to create the SNMP target, which will appear in the **SNMP Target** list. You can modify any of the values entered during the creation process by going to the **Edit Parameters** tab for the selected target (double-click on the target in the SNMP target list).

You can delete an SNMP target by selecting it and then clicking on the **Delete** button above the target list.

MIB

You can download the MIB file with the SNMP trap definitions for Active Circle. This file is a text file containing the SNMP trap definitions used for supervision notes. Go to the **SNMP Target** item in the **Advanced** view of the Administration Tool and click on the **MIB** link:

Figure 30.3. Download Active Circle MIB file



Below is a listing of the contents of the MIB file at the time of publishing:

```
ACTIVECIRCLE-TRAP-MIB DEFINITIONS ::= BEGIN

IMPORTS
    enterprises,
    MODULE-IDENTITY
        FROM SNMPv2-SMI;

activeCircle MODULE-IDENTITY
    LAST-UPDATED "201206251626Z"
    ORGANIZATION ""
    CONTACT-INFO
        ""
    DESCRIPTION
        ""
    REVISION "201206251626Z"
    DESCRIPTION
        "Initial version."
    -- 1.3.6.1.4.1.21399
    ::= { enterprises 21399 }

activeCircleProduct OBJECT IDENTIFIER
    -- 1.3.6.1.4.1.21399.1
    ::= { activeCircle 1 }

activeCircleStandard OBJECT IDENTIFIER
    -- 1.3.6.1.4.1.21399.1.1
    ::= { activeCircleProduct 1 }

-- Traps

activeCircleNotification OBJECT IDENTIFIER
    -- 1.3.6.1.4.1.21399.1.1.2
    ::= { activeCircleStandard 2 }

activeCircleNote NOTIFICATION-TYPE
    -- 1.3.6.1.4.1.21399.1.1.2.1
    OBJECTS { activeCircleNoteName,
               activeCircleNoteDate,
               activeCircleNoteSource,
               activeCircleLevel,
               activeCircleDescription,
               activeCircleNoteCause
            }
    STATUS current
    DESCRIPTION
```

```
"An active circle supervision note."
--#TYPE "Note"
--#SUMMARY "%s - %s on %s (%s, %s)."
--#ARGUMENTS {3, 0, 2, 4, 5}
--#SEVERITY INFORMATIONAL
::= { activeCircleNotification 1 }

activeCircleNoteName OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.1
::= { activeCircleNote 1 }

activeCircleNoteDate OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.2
::= { activeCircleNote 2 }

activeCircleNoteSource OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.3
::= { activeCircleNote 3 }

activeCircleLevel OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.4
::= { activeCircleNote 4 }

activeCircleDescription OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.5
::= { activeCircleNote 5 }

activeCircleNoteCause OBJECT IDENTIFIER
-- 1.3.6.1.4.1.21399.1.1.2.1.6
::= { activeCircleNote 6 }

END
```

Chapter 31. Administration Log

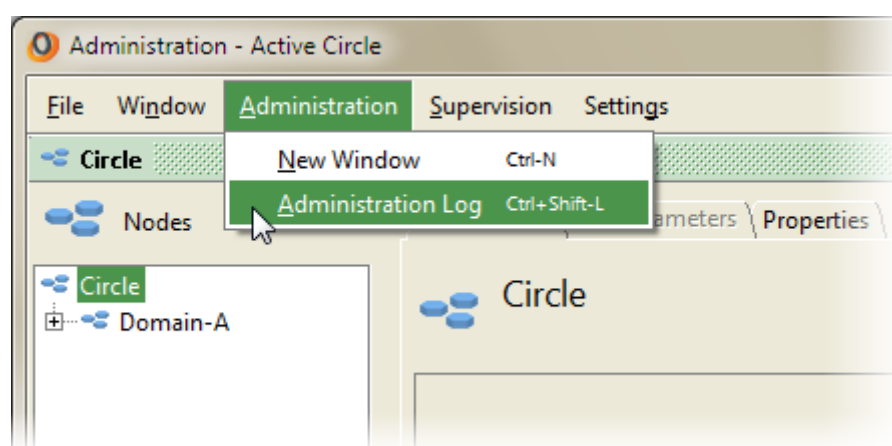
All the actions of administrators, such as the creation and deletion of nodes, shares, users, etc.; system events, such as node start-up and shutdown, the election of a new leader for an HA cluster; as well as the supervision-notes are all stored in the Administration Log.

The system administrator can use the Administration Log tool to search through the events, filtering by node name, by Active Circle module (domain, node, partition, etc.) or by date.

Accessing the Administration Log

The Administration Log has its own window within the Active Circle Administration Tool. To access it, click on the menu **Administration** and select **Administration Log**.

Figure 31.1. Accessing the Administration Log



Tip

You can also use the keyboard shortcut **Ctrl+Shift+L**.

The Administration Log is divided into two tabs:

- **Log**
- **Administration**

Log

The entries are presented in the **Log** tab in a 3-column table listing the **Date** of the event, the **Message**, and the **Node** that generated the message.

Figure 31.2. Log tab

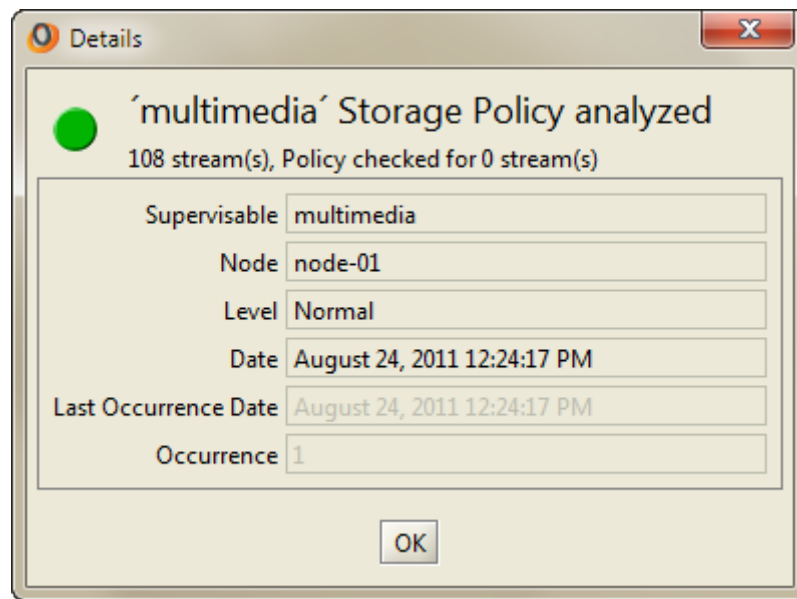
Date ▼	Message	Node
August 24, 2011 12:24:17 PM	● 'multimedia' Storage Policy analyzed	node-01
August 24, 2011 12:24:17 PM	● 'multimedia' Storage Policy analysis started	node-01
August 24, 2011 12:24:17 PM	● 'projects' Storage Policy analyzed	node-01
August 24, 2011 12:24:17 PM	● 'projects' Storage Policy analysis started	node-01
August 24, 2011 12:23:19 PM	● Node 'node-02' started	node-02
August 24, 2011 9:32:04 AM	● Cluster 'Sample Cluster' elected on 'node-01'	node-01
August 24, 2011 9:31:49 AM	● Cluster 'Lock Manager' elected on 'node-01'	node-01
August 24, 2011 9:31:49 AM	● Cluster 'Policy Management' elected on 'node-01'	node-01
August 24, 2011 9:31:07 AM	● Node 'node-01' started	node-01
August 23, 2011 7:06:04 PM	● Node 'node-01' stopped	node-01
August 23, 2011 7:03:34 PM	● Node 'node-02' stopped	node-02
August 23, 2011 5:47:36 PM	● E-mail address 'ana-maria@active.circle.com' created	node-01
August 23, 2011 5:47:02 PM	● E-mail address 'giovanni@active-circle.com' created	node-01
August 23, 2011 5:46:39 PM	● E-mail address 'john@active-circle.com' created	node-01
August 23, 2011 5:46:22 PM	● E-mail address 'marie@active-circle.com' created	node-01
August 23, 2011 2:54:24 PM	● Group 'Tech writers' created	node-01
August 23, 2011 2:54:12 PM	● Group 'Developers' created	node-01
August 23, 2011 2:52:57 PM	● Group 'Administrators' created	node-01
August 23, 2011 2:52:42 PM	● Group 'Administrateurs' deleted	node-01
August 23, 2011 2:52:28 PM	● Group 'Operators' created	node-01
August 23, 2011 11:53:03 AM	● Person 'Billy Bordelais' created	node-01
August 23, 2011 9:43:21 AM	● Synchronized directory "Active Circle Directory"	node-01
August 23, 2011 9:43:14 AM	● Directory 'AD, LDAP (Active Directory, OpenLDAP, Su...	node-01
August 23, 2011 9:35:46 AM	● Directory 'AD / LDAP, LDAP (Active Directory, OpenL...	node-01
August 23, 2011 9:34:46 AM	● Directory synchronization "Active Circle Directory"	node-01
August 23, 2011 9:33:59 AM	● Directories deletion "Active Circle Directory"	node-01

Active Circle retains the last 1,000 messages generated by the system. Messages older than that are transferred into log archives that can be imported back into the **Log** view for consultation. See ["Importing Old Log Files"](#) below for more information.

Message Details

You can view details of any message in two ways: by double-clicking on the message in the table or by right-clicking on the message and selecting **Details**.

Figure 31.3. Message Details dialog box




You can also copy the contents of one or more messages and paste them into another application for analysis, for example a spreadsheet. Because the data is formatted as tab-separated values, it pastes directly into columns.

To copy the message details, click to select the message (or **Ctrl**+click to select multiple messages) you wish to copy. Then right click on the selected messages and click **Copy**.

Filtering the Logs


You can filter the **Logs** view by any combination of the values displayed in the columns: **Date**, **Message**, and **Node**. You can also choose not to display messages that are only informational (identified by a green dot ●).

Figure 31.4. Filters area



Using the **Date** field, you can apply a very precise filter to the view according to a range of dates. To do so, click on the **arrow** in the **Date** field.

Figure 31.5. Filter by Date



Use the **Modify** buttons to change the start or end dates — the **To** or **From** fields respectively. A calendar pop-up allows you to set the date and time for the filter.

Figure 31.6. Select the Date

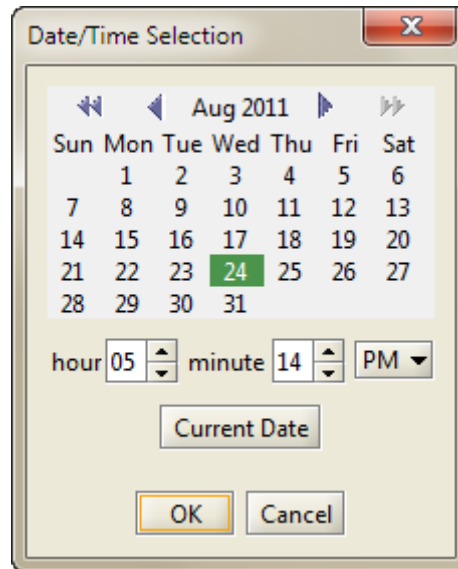
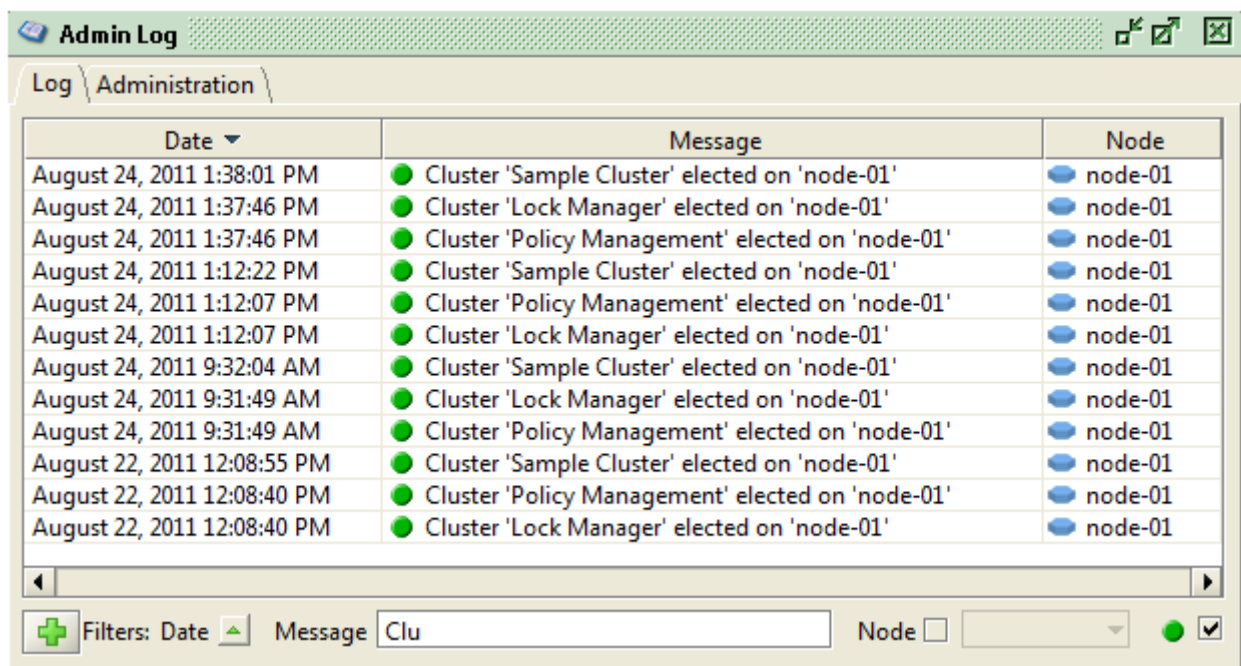


Figure 31.7. Filter by Message



The **Message** field allows you to filter the view by the contents of the message. The log view will be filtered as you type. Any message containing the character sequence you type will be displayed. The field is not case-sensitive.

You can also filter by the node from which the message originated by checking the **Node** box and selecting the appropriate node.



Note

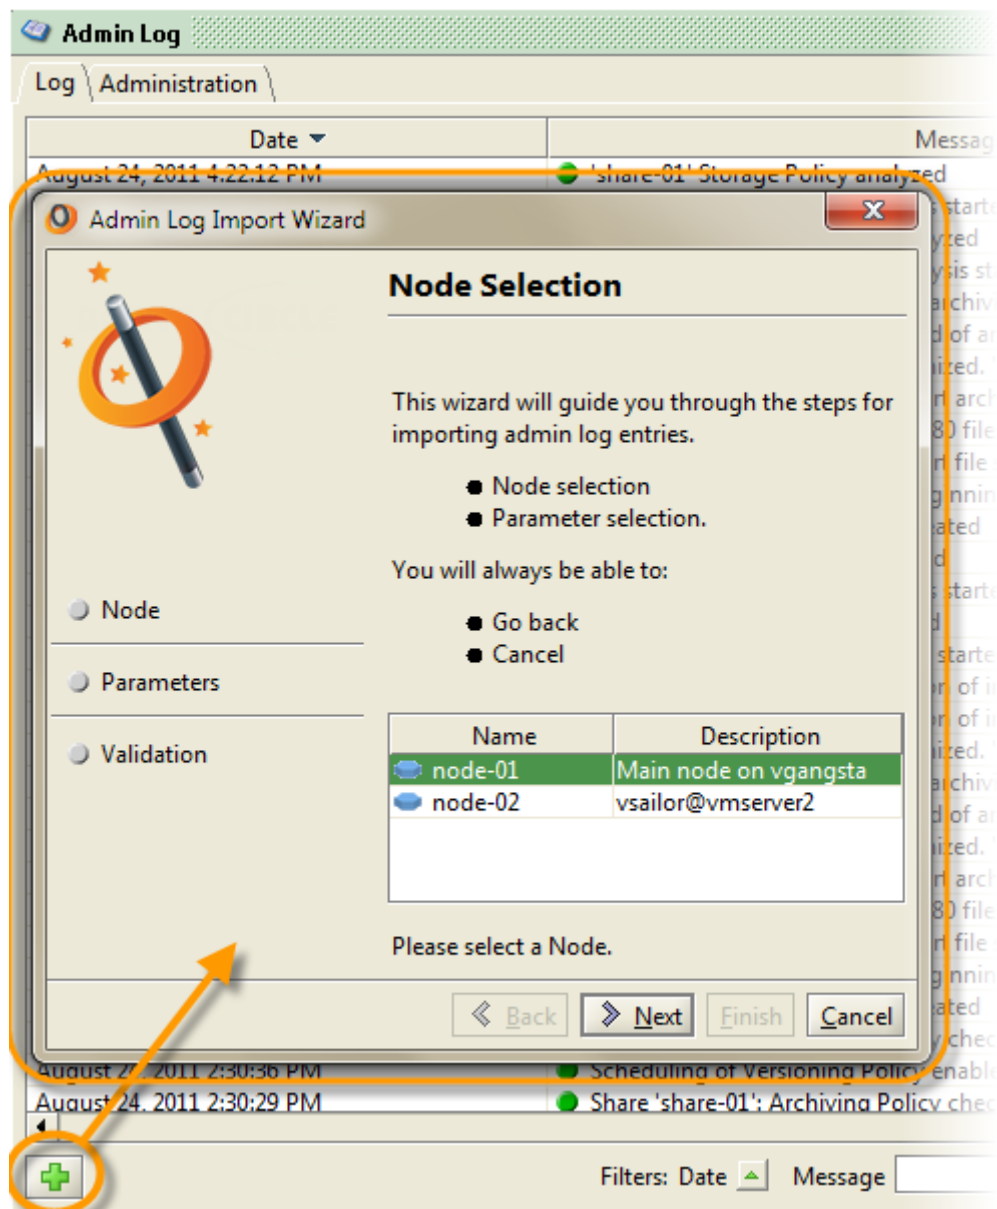
Unchecking the box to the right of the node filter (identified by a green dot ●) removes all informational messages, so that only messages that indicate a problem in the system are displayed.

Importing Old Log Files

The Administration Log displays the last 1,000 messages generated by the system. Above that number, the messages are archived. If you need to consult messages that have been archived, you must therefore import them into the Administration Log from one of the nodes where the archives are stored.

To import the archived administration logs click on the green **plus sign** at the bottom left of the log window.

Figure 31.8. Import Wizard



The Administration Log Import Wizard starts, and it will guide you through the process. You must select a node from which to import the archived log files and then the log files to import. This can be done either by selecting the log files directly or by defining a date range.

Administration

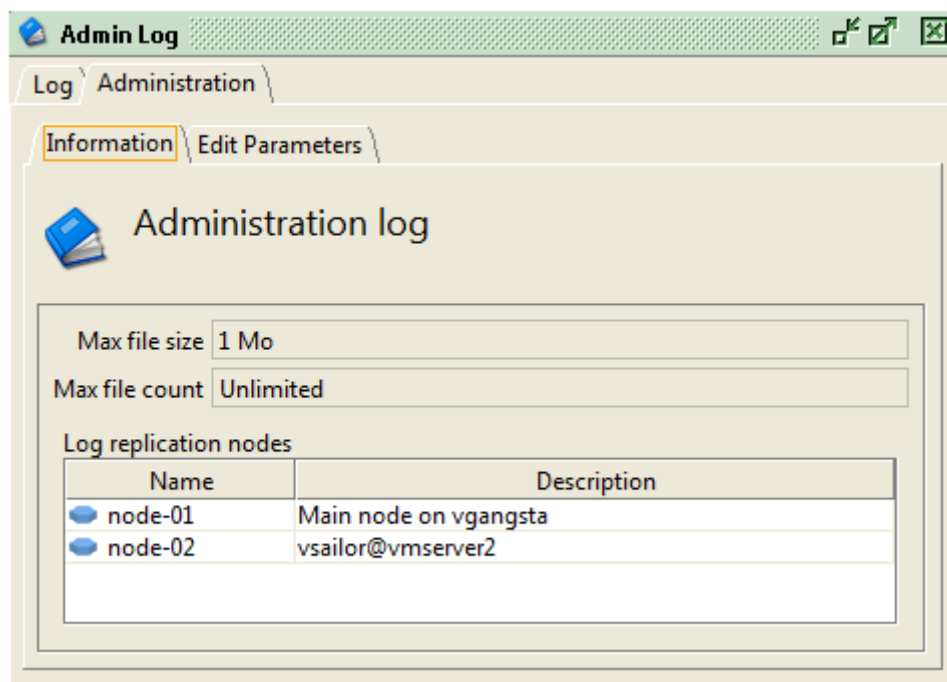
The **Administration** tab is where you can configure the administration logs. It is divided into two sub-tabs:

- [Information](#)
- [Edit Parameters](#)

Information

The **Information** tab presents the administration log parameters that can be modified in the [Edit Parameters](#) tab.

Figure 31.9. Information tab



Max file size

Maximum size allowed for a log file. When a log file reaches this size, the archive is closed and a new one is created. The default value is **1 MB**. The default value can be changed using the [circle property](#) `adminlog.currentSize`.

Max file count

The maximum number of log files to be stored. Once this limit is reached, the system will begin deleting the oldest log files. The default value is **Unlimited**.

Log replication nodes

Nodes to which the archived log files will be copied.

Edit Parameters

You can modify the parameters associated with the administration log in this tab.

Figure 31.10. Edit Parameters tab

The screenshot shows the 'Admin Log' window with the 'Edit Parameters' tab selected. The window has a breadcrumb trail: Log > Administration > Information > Edit Parameters. The main area contains several input fields and a table.

- Name:** Administration log
- Description:** (empty)
- Max file size:** 1 (with a spinner)
- Max file count:** 1 (with a spinner) and a dropdown menu showing 'Unlimited' and 'Limited'.
- Log replication nodes:** A section with 'Add' and 'Delete' buttons and a table.

Name	Description
node-01	Main node on vgangsta
node-02	vsailor@vmserver2

At the bottom right are 'Apply' and 'Cancel' buttons.



Warning

If you set the parameter **Max file count** to a value other than **Unlimited**, the system will delete the archived log files above that number. Use with caution.



Note






If you encounter performance problems due to a large number of administration notes for each share in the administration log, you can reduce the number of allowed log entries per share. For more information, please see [“Share Note Count Limit”](#).

Chapter 32. Supervision

Active Circle provides a notification system in the form of supervision notes. The supervision interface provides the administrator with a detailed overview of the system's performance and notifications in the event of problems. Notifications can be configured at several levels of functionality and for various recipients.

Most system events will produce a supervision note, even normal events. Each message is assigned a color-coded level to help prioritize issues and assist in troubleshooting.

Table 32.1. Supervision Note Color Coding

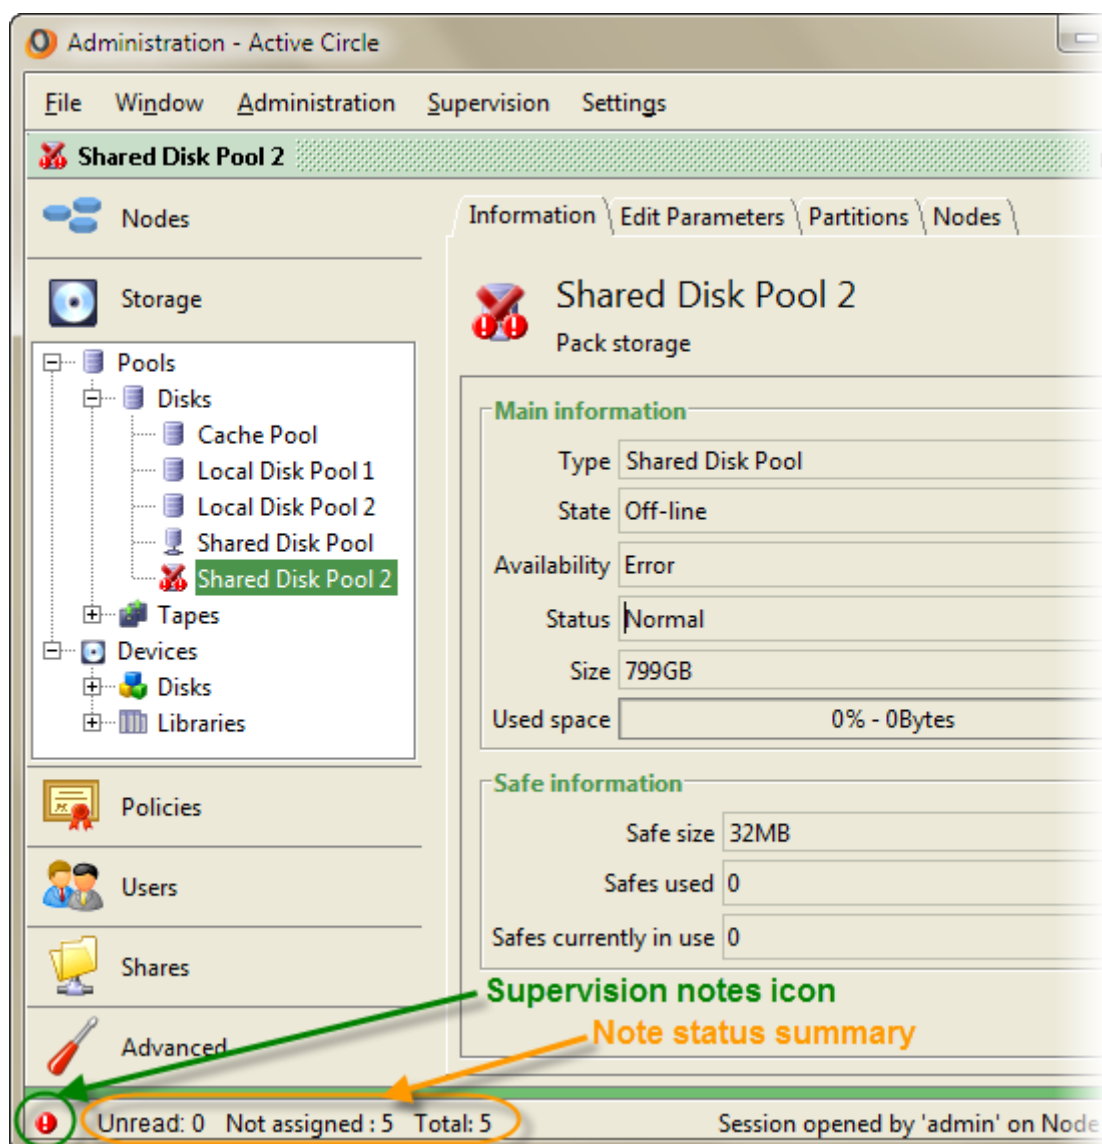
Icon	Meaning
	Critical (red)
	Major (orange)
	Minor (yellow)
	Notice (blue)
	Normal (green - information only)

Where possible, the supervision notes suggest what corrective actions to take in order to resolve the issue that caused the note. When the problem is resolved, the system detects this and closes the corresponding note automatically.

You can access the supervision note window in three different ways:

- by clicking on the menu **Supervision** and selecting **New Window**,
- by using the shortcut **Ctrl+Shift+N**, or
- by clicking on the supervision note icon in the lower left-hand corner of the Administration Tool window.

Figure 32.1. Accessing the supervision note window



Tip

The note icon (located in the lower left-hand corner of the Administration Tool window) displays the color of the highest level warning currently open in the supervision notes, the number of unread notes, the number of notes currently not assigned, and the total number of notes.

The supervision notes window is divided into three tabs:

Note

Where notes are displayed and managed.

Note Registrations

Where notification options are configured.

Advanced

Where you can troubleshoot issues by navigating through the dependencies of devices and services.

Note

The **Note** tab presents all of the supervision notes in a sortable table:

Figure 32.2. Note tab

Name	Source	Node	Status	Level	Creation Date	Last Occurrence
Archive:common#12: archiving failed be...	common - node2	node2	New	Minor	May 24, 2013 12:01:05 AM	May 24, 2013 12:01:05 AM
Partition already allocated	Partition data2 - ...	node1	New	Notice	May 23, 2013 5:50:04 PM	May 23, 2013 5:50:04 PM
Archive:multimedia#10: interrupted archi...	multimedia - no...	node3	New	Major	May 23, 2013 12:01:17 AM	May 23, 2013 12:01:17 AM
Archive:common#9: interrupted archivin...	common - node2	node2	New	Major	May 21, 2013 12:00:11 AM	May 21, 2013 12:00:11 AM
Archive:multimedia#7: archiving failed b...	multimedia - no...	node3	New	Minor	May 20, 2013 12:01:06 AM	May 20, 2013 12:01:06 AM
Archive:common#8: archiving failed bec...	common - node2	node2	New	Minor	May 20, 2013 12:01:05 AM	May 20, 2013 12:01:05 AM

☐ Display The Closed Notes

Archive:common#12: archiving failed because the pool 'TAR 1' is Off-line
 The current state for the pool 'TAR 1' (Off-line) is not compatible with an archiving operation. It should be set to On-line.

Date	Operation
May 24, 2013 12:01:05 AM	Note created
May 27, 2013 10:41:16 AM	Note read by admin

Name

The name of the event that caused the note creation. The name is preceded by a note icon to indicate its level.

Source

The system object that generated the note.

State

The current state of the note. Possible values are:

- **New:** The note has not yet been assigned to a user for resolution.
- **Assigned:** The note has been assigned to a user.
- **Pending:** The note was assigned to a user, but that user then rejected the assignment. In order for it to be resolved, a user must accept responsibility for the note.
- **Closed:** Either the user who accepted the note has resolved the issue and closed the corresponding note or the system has closed the note automatically when the cause of the problem was removed.

Level

The level of importance of the note. For possible values, please see [Table 32.1, "Supervision Note Color Coding"](#)

Creation Date

The date and time the note was generated.

Last Occurrence Date

The last time the event occurred.

Occurrence

The number of times the same note has been generated.

Responsible

The user who has accepted responsibility for resolving the issue.

Cause

The reason the note was generated (if available).


Filtering the Notes

There are several options for filtering the view of the supervision notes.

By default, the view filters out closed notes. These can be displayed by checking the box **Display the Closed Notes**. You can hide them again by unchecking the box.

You can also perform a filtered search for notes by clicking on the **Filter** drop-down menu (identified by a magnifying glass) and selecting one of the search fields.

Figure 32.3. Filter Fields



For example, if you were to select **Source** as the search filter and then entered *Disk Pool 1* in the search field, you would see only notes relating to that object — in this case, a Pool — displayed in the notes view.

You can also filter the view by the level of importance of the supervision note by using the check boxes corresponding to the levels. By default, all levels are displayed. Unchecking a box hides notes corresponding to that level.

Supervision Note History

When you click on a supervision note, its history is displayed in the lower half of the **Note** tab (see [Figure 32.2, “Note tab”](#)). Whenever a user performs an action relating to a note, there is an option of providing a description of the action, information which may prove useful in case the event happens again.

Working with Notes

When a note is created that requires user intervention, the administrator responsible for the action must accept responsibility for it in the supervision note view, by selecting the note and clicking the **Accept** button. A dialog box is displayed where the user can type a description and, optionally, choose to close the note. When the user clicks **OK**, the user name will appear next to the note in the **Responsible** column.

When an accepted note is selected, the label on the **Accept** button is changed to **Reject**. The user can undo the accept action by clicking **Reject** (this can also be done by right-clicking on the note). The note's state is then switched to **Pending** and it can then be accepted by another administrator.

When the problem causing note has been solved, the user can move the note to the **Closed** state, either by clicking the **Close** button or by right-clicking on the note and selecting **Close**.

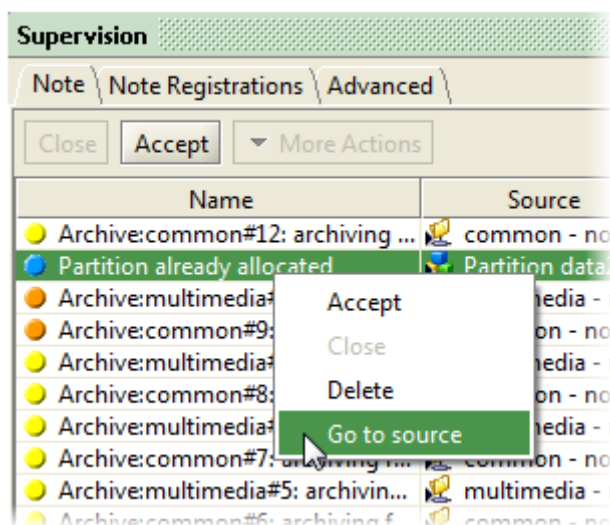
Finally, a user can click the **Delete** button (or right-click and select **Delete**) to remove a note from their personal view. This does not delete the note from the system, but removes it from the current user's view.

Source Details

Supervision notes facilitate the troubleshooting process by providing direct access to the dependency tree. Right-click on the note and select **Go to source**. This will open the

Advanced tab at the position of the source object in the system object tree, allowing you to view the processes on which it depends and those process that depend on it.

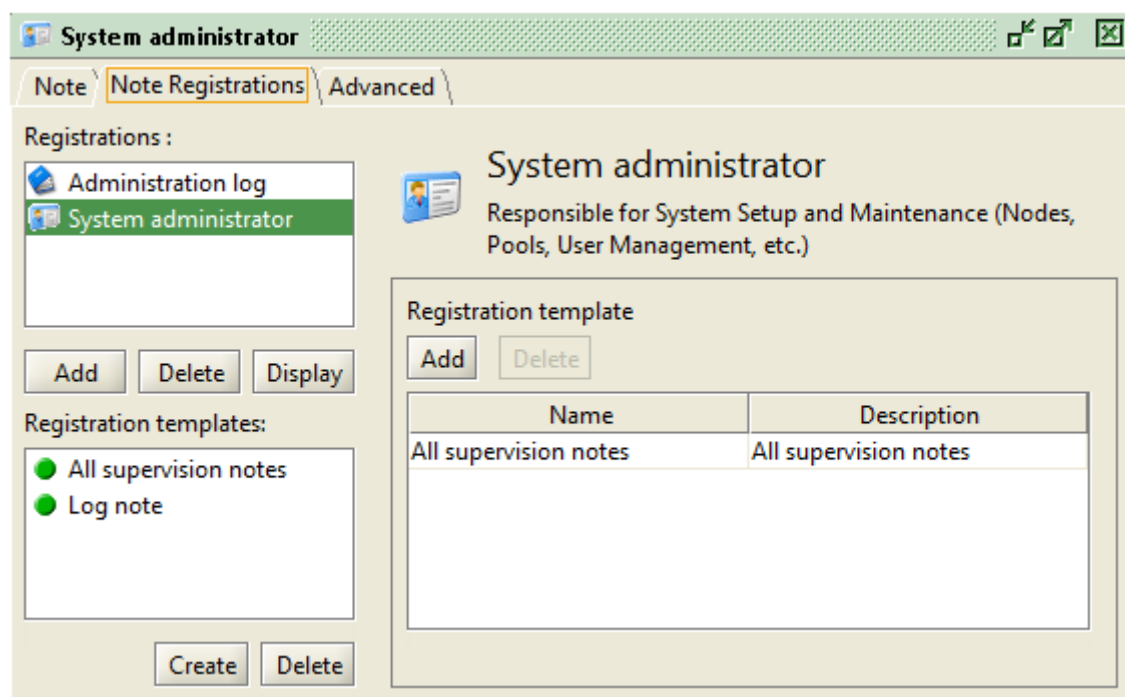
Figure 32.4. Source



Note Registrations

The **Note Registrations** tab allows you to manage how system event messages are communicated to either users or to other systems.

Figure 32.5. Note Registrations tab



Access to the supervision notes is provided by a system of subscriptions, called **Registrations**. The subscribers are the recipients of the notifications, and a **Registration Template** defines what type of information they will receive.

The Active Circle administrative user, `admin`, automatically receives notifications of all system events. In an environment with only one or a few administrators, this default registration may be sufficient. To add more subscribers, use the **Note Registrations** to create a fine-grained system of alerts whereby different users or groups of users are notified of the system events that concern them. Such an approach is practical in a system where there are several people involved in the maintenance of the system, each with clearly defined roles.

When creating a registration, you must consider who (or what, in the case of an [SNMP target](#)) will receive the notification and what information they will receive.

You can choose to send notifications to five different types of subscriber:

- User
- Group
- Profile
- E-mail Address
- SNMP Target



Note

Although there are separate types for `User` and `E-mail address`, it is possible to provide e-mail notification when subscribing users.

If you intend to provide e-mail notification to users, those users must have e-mail addresses associated with their user accounts (for more information, see [“E-Mail Addresses”](#)) and Active Circle must be configured to send e-mails (for more information, see [Chapter 29, Outgoing SMTP Server](#)).

Registration Templates

Before creating a registration, you must first consider which template you intend to use. There are two default templates:

All supervision notes

Which provides access to the supervision notes

Log notes

Which provides access to the logs

To create a new **Registration template**, click on the [Create](#) button in the **Registration templates** area of the **Note Registrations** tab. The **Create Registration Template** dialog box appears.

Figure 32.6. Create Registration Template dialog box

Name	A name to identify the Registration template
Description	A description to help identify the template
Note type choice	The type of note that should be sent. The choices are: <ul style="list-style-type: none"> • Supervision notes • Device notes • Directory notes • Tape notes • Quota notes
Minimum note level	The minimum importance level of the note, that is all notes of this level and higher. The choices are: <ul style="list-style-type: none"> • Normal • Notice • Minor • Major • Critical
Limit to a specific node	By activating this check box, the registration can be limited to notes originating from the selected node.
Limit to a specific domain	By activating this check box, the registration can be limited to notes originating from the selected Active Circle domain.
Limit to a specific source	By activating this check box, the registration can be limited to notes originating from the selected Active Circle share.

Click **OK** to validate the choices. The **Registration template** is created, added to the template list and is ready to use in the context of a **Note Registration**.

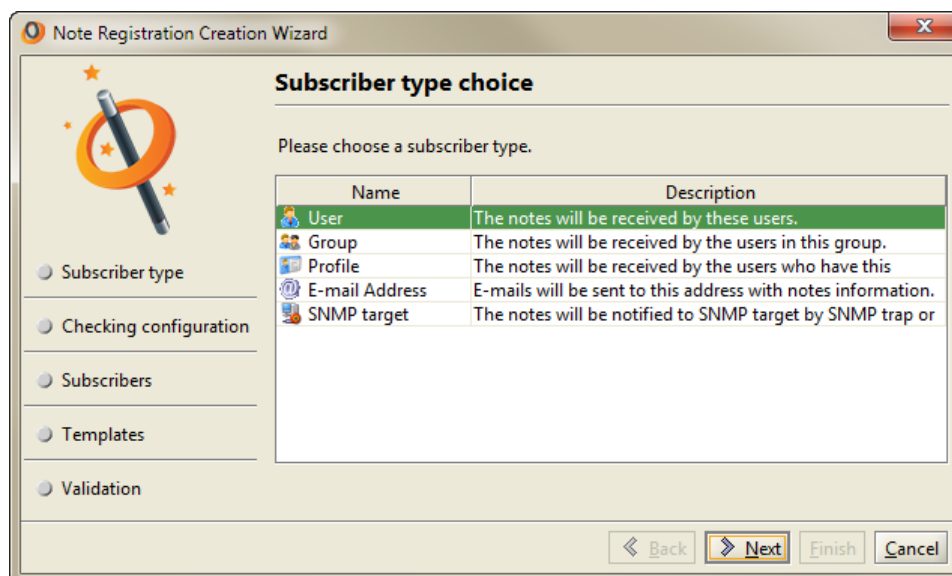
Registrations

The **Registrations** provide users access to the supervision notes.

The Active Circle `admin` user is automatically subscribed to the notes and requires no additional configuration. If you want another user to be able to consult the notes, this user must have an Operator or Administrator profile that allows them to access the notes, and they must be subscribed to the notes.

To create a new registration, click on the **Add** button below the list of registrations **Registrations**. The **Note Registration Creation Wizard** opens to its introductory screen. Click the **Next** button.

Figure 32.7. Note Registration Creation Wizard



Select the subscriber type:

- User
- Group
- Profile
- E-mail Address
- SNMP target

The next step in the wizard depends on your choice of subscriber type. Select the subscriber type you want to configure and click on **Next** to continue. The **Subscribers choice** window appears.



Note

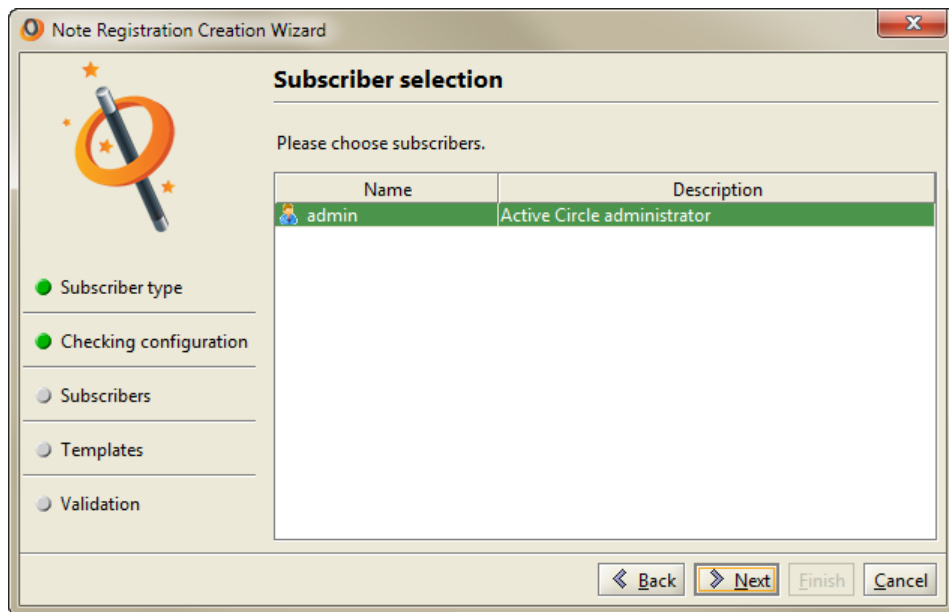
If all elements of the selected subscriber type have already been subscribed, you will receive the following information message:

All elements of this type are already subscribed. Please create a new one or select another type.

For example, if you selected an e-mail address or an SNMP target, this means that all the addresses or targets defined in the system have already been subscribed (or that no e-mail addresses or SNMP targets have been defined).

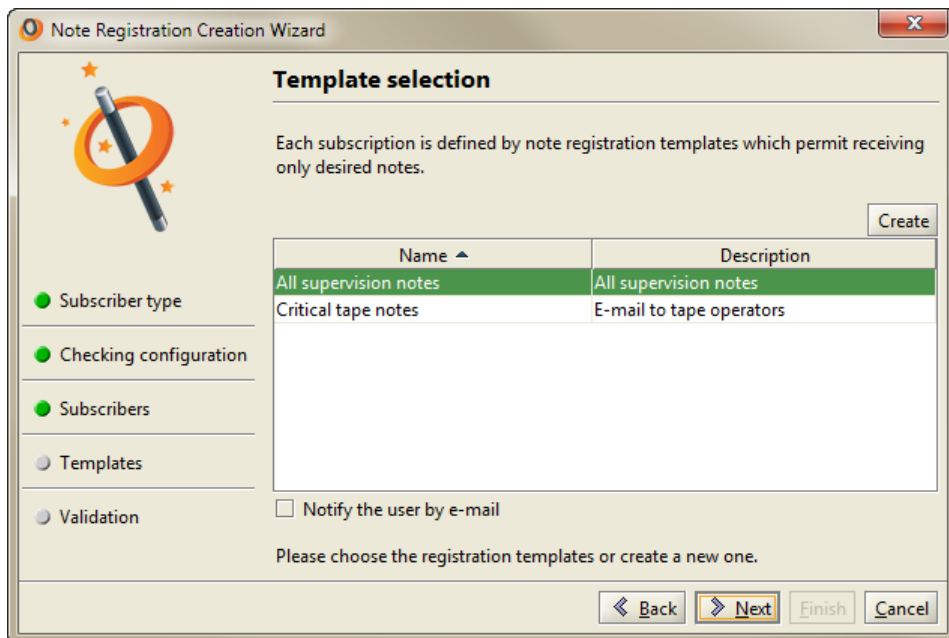
You must [add an e-mail address](#) or [create an SNMP Target](#) before you can add a note subscription for it.

Figure 32.8. Note Registration Creation Wizard - Choose subscriber



In this window you must select the particular subscriber to receive the notification. This could be a user name, an e-mail address, an SNMP Target, or other, depending on your choice. Once you have made your selection, click **Next** to continue.

Figure 32.9. Note Registration Creation Wizard - Choose template



Choose the template which determines what level of notes will be sent. If the template you need is not created yet, you can create it from this window by clicking on the **Create** button. Click **Next** to continue.

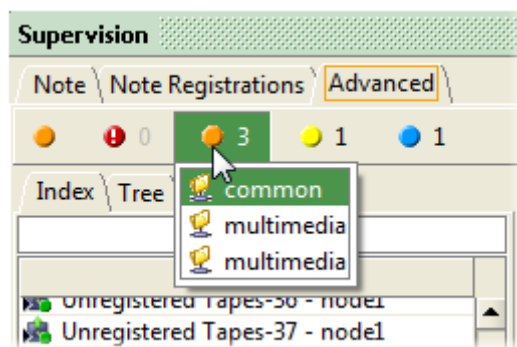
The final screen of the wizard allows you to review your choices before creating the registration. If the registration is correctly configured, click **Finish** to create the registration.

Advanced

The **Advanced** tab in the supervision notes is a troubleshooting tool in which you can view the current state of any system component and navigate through its various interdependencies. For certain objects you can also stop and start individual services. If a note is raised by a particular component, you can click on that component in the **Advanced** tab to view which components it depends on and which depend on it.

At the top of the tab, there is a row of buttons indicating the number of system objects that have supervision notes of a certain level associated with them. You can click on a level button to display the list of objects:

Figure 32.10. Advanced tab – Note level buttons



Click on an object in the list to go directly to the advanced information window for that item.

There are two sub-tabs providing different methods for viewing the components:

- Tree** Displays the information in a navigable tree structure.
- Index** Displays the information in a flat, list view arranged alphabetically.

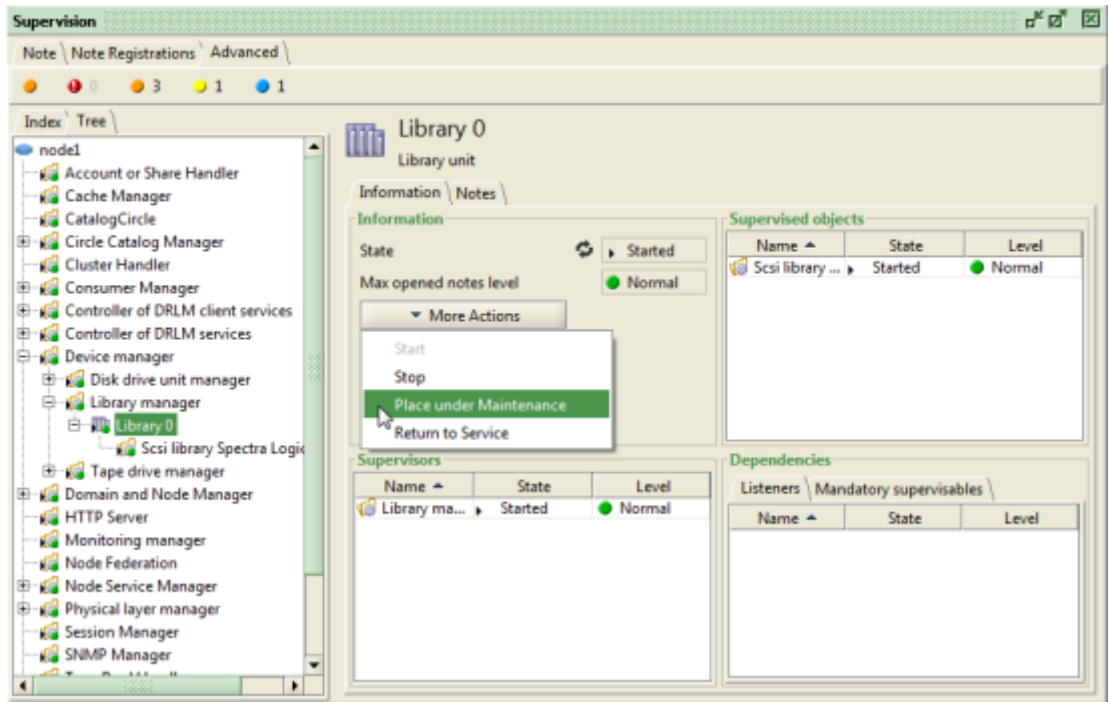
The two views present the Active Circle system components in different ways, but the information displayed about those components and their interdependencies is the same. The views are linked: if you display a component in the **Index** view and then click on the **Tree** view, information on the same component will still be displayed.



Tip

You can access the dependency view in the **Advanced** tab directly from a supervision note by right-clicking on the note and selecting **Go to source**.

Figure 32.11. Advanced tab Tree view



Clicking on a component displays the **Information** tab, which shows its current state, whether or not there are any open supervision notes related to the component, its **Supervisors**, its **Supervised objects** (if any), and its **Dependencies**.

The **Information** pane also contains a button named **More Actions**. For certain types of objects and notes, there are additional actions available directly from this supervision view when you click on this button (starting and stopping shares or pools, placing under maintenance, etc.).

Use the **Notes** tab for the component to display a list of all the open supervision notes associated with it.

Share Note Count Limit

If large numbers of supervision notes per share are stored on a node, it can impact the node performance. In such a situation, you might want to reduce the maximum number of notes that can be stored for each share.

The default limit for the number of entries per share in the Administration log and the supervision notes is 512. You can change this value by modifying the `supervision.noteCountMax` circle property. The value must be between 10 and 512. For more information on editing circle properties, see [“Circle Properties”](#).

Appendix A. Local Variables

The behavior of an Active Circle installation is determined by the configuration of its Nodes. This configuration information is stored on each Node in the file `/activecircle/.vars`.

Most variables are accessible through the **Properties** tab for the Circle Node elements (see for example [“Circle Properties”](#) or [“Properties”](#)).

Some variables, however, can only be changed by adding a file named `.localvars` to the directory `/activecircle`.



Important

Any change to the local variables in the `.localvars` file requires a restart of the Active Circle service on the Node in order for the changes to take effect. See [Chapter 3, The Active Circle Service](#) or [“Restarting a Node”](#) for information on how to restart the service.

You will find a sample file named `.localvars.sample` in the directory `/activecircle`. This file contains some of the more commonly used variables. The default value for all of these variables is `false`.

Table A.1. Local Variables

Option	Usage
<code>CELL_PROBE_SCSI</code>	When <code>true</code> (default), activates detection of SCSI devices.
<code>CELL_PROBE_SCSI_DISK_DRIVES</code>	When <code>true</code> , activates identification of SCSI disk drives. Only for RAW mode.
<code>CELL_PROBE_SCSI_TAPE_DRIVES</code>	When <code>true</code> (default), activates identification of SCSI tape drives.
<code>CELL_PROBE_SCSI_LIBRARIES</code>	When <code>true</code> (default), activates identification of SCSI tape libraries.
<code>CELL_NO_COMPRESS</code>	When <code>true</code> , deactivates data compression.
<code>CELL_CYPHER</code>	When <code>true</code> , activates data encryption.
<code>SERVICE_INTERFACES</code> <code>NETWORK_INTERFACES</code>	For example, "eth0". The two variables should normally define the same interface.
<code>TCP_PORT_RANGE</code>	For example, "31000-31199,29800,30500-30550"
<code>CLIENT_GUI_ADMIN_PORT_RANGE</code>	Port range used by Administration application. Minimum 20 ports. For example, "31200-31219"
<code>CLIENT_GUI_EXPLO_PORT_RANGE</code>	Port range used by Explorer application. Minimum 20 ports. For example, "31300-31319"

Appendix A. Local Variables

Option	Usage
SHORT_REQUEST_TIMEOUT	Allows you to define the length of the short remote request timeout (only when debug is activated). For example, 36000000.
SERVICE_STOP_TIMEOUT	Allows you to define the length of time before sending <code>kill -9</code> when stopping the Java virtual machine. For example, "120"
CLIENT_GUI_ADMIN_MEM	Define the memory allotted to the Administration Tool client (in MB). The default value is 512.
CLIENT_GUI_EXPLO_MEM	Define the memory allotted to the File Explorer GUI client (in MB). The default value is 256.
CLIENT_COMMAND_MODE_MEM	Define the memory allotted to the command-line client (in MB). The default value is 512.
CLIENT_GUI_VERBOSE	Allows you to switch the GUI clients to verbose mode.
CELL_LOCAL_OPT	This option allows you to define additional properties to modify the behavior of the local node and is generally used by Active Circle support personnel.



Important

Any changes to the local variables should be made in the `.localvars` file, which is retained during the software upgrade process.

Tape libraries

In order to disable detection of attached tape libraries, the following lines in `.localvars` need to be uncommented (the system default value is "true"):

```
setenv CELL_PROBE_SCSI false
setenv CELL_PROBE_SCSI_TAPE_DRIVES false
setenv CELL_PROBE_SCSI_LIBRARIES false
```

Network Interfaces

By default, Active Circle uses all the available network interfaces for client access to the file servers on the nodes. When the nodes have multiple network interfaces, it can be useful to manage the network traffic between the nodes separately from the traffic between file servers and clients. The data between nodes include pack transfer, administration information, synchronization information, communication with the GUI applications, etc. This can be

managed using the following variables by assigning them to the appropriate network device ID:

```
setenv SERVICE_INTERFACES "eth0"
setenv NETWORK_INTERFACES "eth0"
```



Important

- Except for in special cases/configurations, the same network interface should be defined for both variables.
- The clients must have access to the network interfaces defined by these two variables.

Although both variables should define the same network interface, each one represents a different type of inter-node network traffic:

SERVICE_INTERFACES This option is used to define the network interfaces that will be used for communication between nodes and between the nodes and the client applications: the Administration Tool, the File Explorer, and the command line.

NETWORK_INTERFACES This option is used to define the network interfaces that will be used for transferring data between nodes (for example during archiving or data replication) as well as restoring data via the Active Circle File Explorer or the command line.

Encryption

It is possible to encrypt the data stored in Active Circle by adding the following line to the `/activecircle/.localvars` configuration file:

```
setenv CELL_CYPHER true
```



Important

If you choose to use encryption, you must set this parameter on all of the nodes in your Circle.

The Active Circle system employs its own proprietary symmetric encryption algorithm using a single internal key for all configurations.

Port Ranges

You can define a range of TCP network ports for inter-node communication. This can be particularly useful when implementing a Circle in an environment where the network traffic has to pass through a firewall.

The range should include at least 50 ports and must be defined in the `.localvars` file on each node in the Circle. Uncomment the line for `TCP_PORT_RANGE` and specify the port numbers/ranges to use. The below is an example showing the possible syntax for specifying TCP ports:

```
setenv TCP_PORT_RANGE "31000-31199,29800,30500-30550"
```



Note

Active Circle has published a separate Technical Brief titled *Network port and routing requirements* which describes in detail how to define ports in an Active Circle environment. Please contact customer support if you want to obtain this document.

Multipath verification

The system checks all the paths for multipath devices, and by default the device is considered unstable if less than one percent of the paths are available or more than 1000 paths are unavailable. This would provoke a switch of the shared pool cluster to another server. These limits are by default extreme values to prevent unnecessary switching of leader node as long as there are paths in service. These values can be changed using local variables, to adapt the behaviour to your environment.

Uncomment the following lines in the `/activecircle/.localvars` configuration file and modify the values. The default values are the numbers indicated after the variables below:

```
setenv MIN_PERCENTAGE_REACHEABLE_PATH 1
setenv UNREACHEABLE_PATH_THRESHOLD 1000
```

Appendix B. Configuring and Testing Drives and Libraries

This appendix contains additional information on the configuration of Nodes managing tape libraries.



Important

Because of differences in data flow rate, it is not possible for clients to read data directly from tape. Consequently, the system needs a buffer zone on disk to serve as a temporary storage area. This buffer should ideally be a Cache Pool. For information on how to create a Cache Pool, please see “[Cache Pools](#)”.

Configuring the SCSI Bus

You must install and configure the correct drivers for the SCSI bus on a Node to which a tape library is attached. The examples given below are for Nodes running a RedHat/CentOS distribution of GNU/Linux with a minimum kernel version of 2.6.9.x.

Procedure B.1. Installing the SCSI Bus

1. Power off the server and the library.
2. Connect the SCSI cables, making sure not to forget the terminator plug.
3. Power up the library and wait until it is fully started.
4. Power up the server. During the start up, the BIOS should recognize the library and tape drives.



Note

If possible, avoid connecting the library to the same bus as a SCSI disk, or even on the same SCSI card.

Testing the Hardware Outside of Active Circle

In order to verify that the tape library is functioning correctly, it is recommended to have the `mtx`¹ package installed.

Procedure B.2. Test the Hardware

1. Verify in `/proc/scsi/scsi` that the library and the drives are present on the bus and recognized by the operating system.
2. Verify in `/dev` that the tape drives are present in the file system. For example, a library with two drives should have two entries: `/dev/st0` and `/dev/st1`.
3. Load the SCSI Generic driver with the command:

`modprobe sg`

¹ SCSI Media Changer and Backup Device Control. <http://sourceforge.net/projects/mtx/>

Appendix B. Configuring and Testing Drives and Libraries

4. Check that there are entries corresponding to the library and drives in `/var/log/messages`. You should see something like the following:

```
May 5 08:56:59 cell-02 kernel: Attached scsi generic sg0
at scsi1, channel 0, id 1, lun 0, type 8
May 5 08:56:59 cellule2 kernel: Attached scsi generic sg1
at scsi1, channel 0, id 2, lun 0, type 1
May 5 08:56:59 cellule2 kernel: Attached scsi generic sg2
at scsi1, channel 0, id 4, lun 0, type 1
```



Note

"type 8" corresponds to a library, "type 1" corresponds to a drive.

5. Test the drives with the commands:

`mt -f /dev/st0 status`

`mt -f /dev/st1 status`

Should these commands fail, rerun the procedure from step [Step 1](#). If the problem persists, there may be a defect with the cables, the tape drives, or the SCSI card.

6. Insert the tape cartridges into the library, making sure the bar codes are correctly attached.

Procedure B.3. Configuring SCSI Target IDs

Go to the configuration menu for the SCSI card and do the following:

1. Select the ID for the library.
2. Select the IDs of the tape drives located in the library.
3. Assign the IDs to the drives in the same order as the tape drive slot addresses in the library. For example, if the drive slots have the addresses "480" and "481", the drive in slot 480 must have an ID lower than the drive in slot 481.

In some library configuration tools, the slot addresses in a library are referred to as "SCSI addresses". These addresses are unrelated to the addresses on the SCSI bus and are only related to the library, where they are used for moving the picker and taking inventory.

A SCSI bus address is made up of three whole numbers from between 0 and 15: channel number, SCSI target ID, and the LUN (logical unit number). A slot address is a single number the format of which varies among manufacturers and library models.

Appendix C. Configuring external software for use with HA Clusters

The information provided here is intended to help the administrator of an Active Circle system to configure third-party software to work in conjunction with Active Circle. This information cannot, however, replace the documentation provided by the creators of the software and it is strongly recommended that you consult the original documentation.

For information on configuring DNS servers, please go to [the BIND9 website](#).

For information on configuring OpenVPN, please go to [the OpenVPN website](#).

Configuring DNS Clusters

A DNS Cluster dynamically associates the lead node of a HA Cluster with a DNS host name. To implement this, you must be using BIND9 as your DNS server.

On the server, enter:

```
dnssec-keygen -a HMAC-MD5 -b 512 -n USER clustername.domain.com
```

where `clustername.domain.com` is the key name (replace with values corresponding to your configuration).

An identical key is written to two files with the extensions `.key` and `.private`. The names of both of the files will begin with the letter K followed by the key name and some additional generated text.

The contents of the `.private` file will look something like this:

```
Private-key-format: v1.2
Algorithm: 157 (HMAC_MD5)
Key: Yr+TkR+mHQ2ZAcD5KEMEWB1DM/QJ+IBOgCT2qCbWA==
```

The key must be declared in the BIND configuration, either directly or through an include statement.

```
key clustername.domain.com. {
algorithm HMAC-MD5;
secret "Yr+TkR+mHQ2ZAcD5KEMEWB1DM/QJ+IBOgCT2qCbWA==" ;
};
```

You must then allow dynamic updating of the DNS with this key:

```
zone "domain.com" {
type master;
file "domain.com";
allow-update { key clustername.domain.com.; };
};
```

Configuring Scripted Clusters with OpenVPN

In order to implement a scripted cluster, OpenVPN must be installed on each node in the HA Cluster as well as on any client machines that wish to access the cluster via this connection.

Installing OpenVPN

This section gives an example of the process for installing OpenVPN using RPM packages. The minimum version of OpenVPN required is version 2.0.

OpenVPN must be installed on all the nodes in the cluster (server installation), as well as on the client machines.

You can obtain pre-packaged versions of OpenVPN for several GNU/Linux distributions from the operating system repositories or at [the OpenVPN website](#).

In most cases, this will also require the installation of the LZO package.

Example of installation commands on CentOS:

```
yum install lzo
```

```
yum install openvpn
```

To ensure that the OpenVPN service is not executed automatically, on the servers run the command:

```
chkconfig --list openvpn
```

If the service is activated for certain runlevels, run the command:

```
chkconfig --del openvpn
```

Configuring OpenVPN

OpenVPN clusters can be configured using the Active Circle command **accluster**. For more information about this, please see the *Active Circle Command Line Guide*.

Running OpenVPN

You can manage OpenVPN on a node through the use of a shell script that allow you to start, stop, and restart OpenVPN.



Note

If the **service** command is not available for your distribution, you will have to adapt the scripts to your environment.

The script must be located in a specific directory relative to the installation directory: `cell/cluster/scripts`.

Appendix D. Managing Windows ACLs

Active Circle Shares support the security attributes of Microsoft Windows Access Control Lists (or ACLs). This means that you can view and configure folder security in exactly the same way as you would with a traditional Windows share.

There are two important points to keep in mind when managing Windows ACLs (permissions) on Active Circle Shares:

- To avoid any potential problems linked to being denied access, it is recommended that the person performing the operation be a domain administrator.
- The same access rights should be set at the root of the source and destination folders in order to ensure that the access rights are correctly inherited.

Any modifications to the ACLs made through Windows are taken into account when users try to access files, regardless of the protocols used.



Important

ACLs only control access to the *files* and *folders*. How to provide user access to the *shares* is covered under [“Access Controls”](#) in [Chapter 23, Shares](#).



Note

It is possible to transfer files and folders with their ACLs. For this purpose you will need to use a tool that provides these capabilities, such as RoboCopy or Xcopy. For more information, see [Appendix E, Migrating Files to Shares](#).

Editing ACLs

When a Share is first created, the system grants "Full Control" to the user `Everyone` (SID `S-1-1-0`) in the permissions on the files and folders under the share (this user has the name `Anyone` in the Active Circle user interface).

Users and administrators of CIFS shares can view and edit access privileges in Windows Explorer. As is the case when editing access rights for native Windows shares, editing access rights in an Active Circle Share requires that the client machine be part of a domain. A domain controller is therefore required. It is not a requirement that the users be part of this domain.

Windows supports several special users, but Active Circle only manages the users `Everyone` and `CREATOR OWNER`.



Tip

You can also manage file and folder security using the Active Circle File Explorer application or with the **acacl** command on the Active Circle command line interface. For more information, please see the *Active Circle File Explorer Guide* or the *Active Circle Command Line Guide*.

Appendix E. Migrating Files to Shares

The third-party tools RoboCopy and XCopy from Microsoft can be used for transferring a Windows folder structure to or from an Active Circle share using the CIFS protocol. While the recommended approach is to copy *without retaining existing ACLs* (NTFS file permissions), it is possible to use these tools to copy with the Windows ACLs. This section describes both approaches and how to achieve this with each tool. These are the topics covered:

- [Tool overview](#) (Robocopy and Xcopy)
- [Transferring files without Windows ACLs](#)
- [Synchronizing a Folder Structure](#)
- [Transferring files with the existing Windows ACLs](#)
- [Xcopy syntax](#) (alternative procedure)

Tools

Robocopy

Active Circle recommends using Robocopy to copy or synchronize folder structures to shares. This is a very robust file copy tool which is perfectly suited for copying large directory structures with all attributes intact. Robocopy can tolerate interruption to the network services and resume copying on its own. It also has a mirroring mode for keeping file trees synchronized. It is available in different versions in several Microsoft Windows products.

Possible versions of Robocopy to use with Active Circle:

- **XP010** (minimum version) – Part of the [Windows Server 2003 Resource Kit Tools](#)
- **XP026** – Part of the [graphical user interface version of Robocopy \(Microsoft TechNet\)](#)
- **XP027** – Included in Windows Vista, Windows Server 2008 and Windows 7

There is an article on the Robocopy tool available at [Wikipedia](#), which includes download links. For general help on using Robocopy, type **ROBOCOPY /?** on the command line or consult the online Robocopy command reference on [Microsoft TechNet](#). The specific options for using Robocopy in an Active Circle environment are described in the sections below.

Xcopy

Xcopy, which is included in all versions of Microsoft Windows, can be used for copying Windows folder structures with and without ACLs. It is only recommended for smaller copying jobs without advanced attributes or permissions, in environments where network interruptions are not an issue.

There is a knowledgebase article about copying on the [Microsoft website](#). You can also get help by typing **xcopy /?** on the Windows command line. The specific options for using Xcopy in an Active Circle environment are described in the sections below.



Important

To have sufficient privileges to run any of these commands with the necessary options, you need to be running the terminal session where you execute the command as an Administrator.

Transferring files without Windows ACLs

The user account performing the copy operation must have read permission on all the files contained in the folder structure to be copied. The file transfer should be performed using a CIFS connection.

Command syntax (without ACLs)

The following is the specific command syntax to use for copying a folder structure to an Active Circle share *without keeping* the Windows ACLs (permissions) on the files and folders.

```
robocopy {source} {destination} /E /ZB /V /XA:S /A-:ANT /FFT [/XF ~* *.tmp]
```

where:

source	Is a folder specified by a network path (UNC) in the format \\Server\Share\[Folder] or a local folder path in the format C:\Path\Folder.
destination	Is a folder specified by a network path (UNC) to an Active Circle share in the format \\Server\Share\[Folder].
/E	Copies sub-directories, including those that are empty.
/ZB	Uses Restartable mode. If that fails with an "access denied" error, switches automatically to Backup mode (requires backup operator or full control permission).
/V	Produces verbose output, and shows all skipped files.
/XA:S	Excludes system files
/A-:ANT	Turns off the attributes that are not supported by Active Circle in the copied files: A – Archive N – Not content indexed T – Temporary
/XF	(Optional). Excludes files with the specified names, paths, or wildcard characters. Add /XD to exclude directories.
/FFT	Uses FAT File Times, meaning 2-second granularity. This is needed to eliminate unnecessary copying when going between and NTFS and Linux/Unix/FAT or emulated file systems. If this switch is not specified, many (or all) files will always be seen as more recent when copying to Active Circle, resulting in unnecessary copying of unchanged files.

Note that the **/COPY** option has been omitted, since the default value **/COPY:DAT** is all that is needed in this scenario (copying without ACLs).

If you want to create a log file for the operation, you can add the command option **/LOG:filename**. This will redirect output to the specified file.

You may also use the same command to copy a folder structure from an Active Circle share to another share. In that case you should reverse the source and destination specifications, use **/Z** instead of **/ZB**, and omit (or change) the **/XF** specification.

Synchronizing a Folder Structure

Using Robocopy's mirroring mode, you can keep directory trees synchronized. This mode is activated by adding the **/MIR** switch to the Robocopy command line. Example:

```
robocopy {source} {destination} /E /ZB /V /XA:S /A-:ANT /FFT /MIR
```

The mirroring mode can be used for synchronizing both with or without ACLs, and to or from Active Circle shares.



Important

For synchronization to work properly with an Active Circle share, it is *imperative* that the **/FFT** switch is specified at the same time, to avoid unnecessary copying of unchanged files.



Warning

Use **/MIR** very carefully. Keep in mind that in the mirroring mode, Robocopy will also delete files. If you specify **/MIR** with an existing destination directory by mistake, Robocopy can quickly delete a large amount of data from the destination.

Transferring files with Windows ACLs

Copying a folder structure to an Active Circle share while retaining the NTFS file security rights (Windows ACLs) is not the recommended method for file migration. Proceed only if you have good knowledge of the implications of this and feel comfortable with these concepts. Consider whether it might be more practical to define the ACLs directly on the destination share instead of migrating these attributes. If the migration of ACLs fails, you will have to reinitialize all the ACLs on the destination and then reapply them. There are some important prerequisites to verify before you start migration using this method.



Note

If you need to reinitialize ACLs, this can be done using the **acacl --reset** command on the Active Circle command line. For more information, please see the *Active Circle Command Line Guide*.

Prerequisites

If the ACLs must be retained during the file migration, the following requirements apply:

- The security identifiers (SID) assigned to the users and groups in the ACLs must be known to Active Circle. This means that the Active Directory of the domain must be imported and synchronized with Active Circle.
- The source and destination folders *must not* inherit their ACLs. The ACLs should be explicitly defined.
- The user performing the copy operation must have full control in the permissions on the data being copied. This user must also have all rights on the Active Circle destination share.

Command syntax (with ACLs)

The following is the specific command syntax to use for copying a folder structure to an Active Circle share while *retaining* the Windows ACLs (permissions) on the files and folders. The file transfer should be performed using a CIFS connection.

```
robocopy {source} {destination} /E /ZB /V /COPY:DATSO /XA:S /A-:ANT /FFT  
[/XF ~* *.tmp]
```

where:

source	Is a folder specified by a network path (UNC) in the format \\Server\Share\[Folder] or a local folder path in the format C:\Path\Folder.
destination	Is a folder specified by a network path (UNC) to an Active Circle share in the format \\Server\Share\[Folder].
/E	Copies sub-directories, including those that are empty.
/ZB	Uses Restartable mode. If that fails with an "access denied" error, switches automatically to Backup mode (requires backup operator or full control permission).
/V	Produces verbose output, and shows all skipped files.
/	Specifies the file properties to be copied. Values used:
COPY:DATSO	D – Data A – Attributes T – Time stamps S – NTFS access control list (ACL) O – Owner information (The value U also exists for this option, but it is not applicable because Active Circle does not store audit information).
/XA:S	Excludes system files
/A-:ANT	Turns off the attributes that are not supported by Active Circle in the copied files: A – Archive N – Not content indexed T – Temporary
/XF	(Optional). Excludes files with the specified names, paths, or wildcard characters. Add /XD to exclude directories.
/FFT	Uses FAT File Times, meaning 2-second granularity. This is needed to eliminate unnecessary copying when going between and NTFS and Linux/Unix/FAT or emulated file systems. If this switch is not specified, many (or all) files will always be seen as more recent when copying to Active Circle, resulting in unnecessary copying of unchanged files.

If you want to create a log file for the operation, you can add the option **/LOG:filename**. This will redirect output to the specified file.

If you want to keep the file trees synchronized, you can use the Robocopy command's *mirroring* mode. The options for this mode are described under [Synchronizing a Windows Folder Structure](#).

Xcopy syntax

Alternatively, you can use the **xcopy** command to copy Windows folder structures.

```
xcopy {source\*} {destination\} /O /E /H /K
```

where:

- * Copies all files under the source directory. Note that if you do not specify the asterisk (*) at the end of the source directory, **xcopy** will overwrite the rights on the *destination* share.
- \ The backslash (\) after the *destination-folder* is necessary to indicate that the target is a directory, and the command will create the directory if it does not exist.
- /O Copies the ACLs and their owners, retaining the existing permissions on each file.
Important: Only specify this if you really want to transfer the Windows ACLs along with the files, and you understand the implications of doing this.
- /E Copies all sub-directories, including empty ones.
- /H Copies hidden files.
- /K Copies file attributes.



Note

This procedure does not override inherited permissions. As with Robocopy, you should make sure ACLs are not inherited before performing this operation.

When you use the Xcopy command with the above options to copy the files, the following inherited permissions apply:

- The ACLs that are directly assigned to the files and folders are retained.
- The ACLs that are inherited from the source parent folder are lost.
- The ACLs of the destination folders are inherited.

Xcopy also has an **/X** option, but it is not applicable because Active Circle does not store audit information.

Circle Properties

This is a list of the [circle-properties](#) that are described in this document.

C

circle properties

- [adminlog.currentSize](#), 308
- [archive.closeTAROnFolderChangeDepth](#), 168, 243
- [archive.disable](#), 112, 114, 162
- [archive.enableExport](#), 114, 169
- [archive.excludedFiles](#), 166
- [archive.externalFilterExcludedFiles](#), 167
- [archive.externalFilterExecutable](#), 167
- [archive.externalFilterPartialIsNok](#), 167
- [archive.fastTreeWalk](#), 168
- [archive.integrityCheckPercent](#), 110
- [archive.keepEmptyArchive](#), 168
- [archive.maxDuration](#), 165
- [archive.maxFilesCount](#), 166
- [archive.maxTarFileSize](#), 166
- [archive.minSize](#), 166
- [archive.triggerOnCacheFilling](#), 244
- [archive.triggerOnCacheFillingNumberThreshold](#), 244
- [archive.triggerOnCacheFillingVolumeThreshold](#), 244
- [archive.validatedSelection](#), 167
- [archive.validatedSelectionExcludedFiles](#), 167
- [archive.validatedSelectionIncludedFiles](#), 167
- [cos.slowPackCopyRatio](#), 152
- [diskPool.autoFailback](#), 88
- [diskPool.clusterSilenceDuration](#), 85
- [media.excludedTapeBarcodes](#), 108
- [media.lifeTime](#), 103
- [media.maxPoolCount](#), 91
- [media.packEnabled](#), 93
- [media.recoveredErrorsThreshold](#), 103
- [media.unusedOfflineTapeCheck](#), 111
- [media.unusedTapeAction](#), 111
- [media.unusedTapeCheckPeriod](#), 111
- [nas.activityFileCompression](#), 272
- [nas.activityFileMaxAggregateLength](#), 272, 273
- [nas.activityFileMaxLength](#), 272
- [nas.activityFileRetentionTime](#), 272
- [nas.audit consolidated](#), 273
- [nas.audit.filteredOutEvents](#), 275

- [nas.audit.maxPendingEvents](#), 272
- [nas.directoryDefaultMode](#), 174
- [nas.fileDefaultMode](#), 174
- [session.log](#), 224
- [session.monitored](#), 223, 273, 273
- [supervision.noteCountMax](#), 321
- [vfs.maxShareCount](#), 231
- [vfs.prefetch](#), 243

Index

A

- access control, [174](#)
 - active circle permissions, [176](#)
 - shares, [174](#)
 - Unix permissions, [174](#)
 - Windows permissions, [175](#)
- access control lists (see ACLs)
- ACLs, [175](#)
 - editing, [175](#)
 - managing, [331](#)
- active circle service
 - extensions, [9](#)
 - starting, [9](#)
 - stopping, [9](#)
- administration log, [303](#)
 - accessing, [303](#)
 - contents, [303](#)
 - edit parameters, [308](#)
 - filtering, [305](#)
 - importing, [307](#)
 - note limit, [309](#)
- administration tool, [13](#)
 - column properties, [23](#)
 - inheriting values, [25](#)
 - loading, [22](#)
 - logging in, [14](#)
 - main window, [17](#)
 - menus, [15](#)
 - settings, [19](#)
 - starting, [13](#)
 - views, [19](#)
- administrator
 - profiles, [172](#)
- archives, [253](#)
- archiving
 - disabling, [162](#)
 - on cache filling, [244](#)
 - prefetching files, [243](#)
 - properties, [161](#)
 - scheduling, [242](#)
 - start incremental, [237](#)
 - verification after writing, [168](#)
- archiving policy, [161](#)
 - constraints, [164](#)
 - create, [162](#)
 - parameters, [165](#)

- auditing, [269](#)
 - activation, [269](#)
 - consolidation, [273](#)
 - enabling on share, [243](#)
 - log file, [270](#)
 - contents, [273](#)
 - location, [271](#)
 - size limit, [271](#)
 - time, [272](#)
- authentication
 - directories, [173](#)
 - protocols, [173](#)
 - users and network, [174](#)
- auto-expand, [98](#)

B

- barcode, [100](#), [104](#)
- blank tapes (pool), [95](#)
- block size, [162](#)

C

- cache, [45](#)
 - clearing, [46](#)
- cache pool, [68](#)
- cell (see node)
- CIFS connection options, [256](#)
- CIFS protocol, [263](#)
- circle
 - definition, [27](#)
 - monitoring nodes, [28](#)
 - node overview, [27](#)
- circle properties, [31](#)
 - adminlog.currentSize, [308](#)
 - archive.closeTAROnFolderChangeDepth, [168](#), [243](#)
 - archive.disable, [112](#), [114](#), [162](#)
 - archive.enableExport, [114](#), [169](#)
 - archive.excludedFiles, [166](#)
 - archive.externalFilterExcludedFiles, [167](#)
 - archive.externalFilterExecutable, [167](#)
 - archive.externalFilterPartialIsNok, [167](#)
 - archive.fastTreeWalk, [168](#)
 - archive.integrityCheckPercent, [110](#)
 - archive.keepEmptyArchive, [168](#)
 - archive.maxDuration, [165](#)
 - archive.maxFilesCount, [166](#)
 - archive.maxTarFileSize, [166](#)
 - archive.minSize, [166](#)
 - archive.triggerOnCacheFilling, [244](#)

- archive.triggerOnCacheFillingNumberThreshold, 244
- archive.triggerOnCacheFillingVolumeThreshold, 244
- archive.validatedSelection, 167
- archive.validatedSelectionExcludedFiles, 167
- archive.validatedSelectionIncludedFiles, 167
- categories, 32
- cos.slowPackCopyRatio, 152
- diskPool.autoFailback, 88
- diskPool.clusterSilenceDuration, 85
- editing, 33
- media.excludedTapeBarcodes, 108
- media.lifeTime, 103
- media.maxPoolCount, 91
- media.packEnabled, 93
- media.recoveredErrorsThreshold, 103
- media.unusedOfflineTapeCheck, 111
- media.unusedTapeAction, 111
- media.unusedTapeCheckPeriod, 111
- nas.activityFileCompression, 272
- nas.activityFileMaxAggregateLength, 272, 273
- nas.activityFileMaxLength, 272
- nas.activityFileRetentionTime, 272
- nas.audit consolidated, 273
- nas.audit.filteredOutEvents, 275
- nas.audit.maxPendingEvents, 272
- nas.directoryDefaultMode, 174
- nas.fileDefaultMode, 174
- session.log, 224
- session.monitored, 223, 273, 273
- supervision.noteCountMax, 321
- vfs.maxShareCount, 231
- vfs.prefetch, 243
- class of service (see policy)
- cleaning tapes (pool), 96
- cloud pool, 117
 - availability, 121
 - create, 117
 - password, 120
 - provider, 118, 118
 - read-only, 121, 121
- cluster (see HA cluster)
- column properties, 23
- connecting to shares, 256
- COS (see policy)
- D**
- defragmentation, 111
- depository (see pool)
- directories, 172, 177
 - authentication, 173
 - configure, 181
 - e-mail, 187
 - external, 179
 - Active Directory/LDAP, 197
 - Active Directory/NT, 205
 - importing, 189
 - NIS, 208
 - OpenLDAP/POSIX, 191
 - merge users, 184
 - order, 181
 - protocols, 172
 - synchronize, 177
 - frequency, 178
 - types, 180
 - users, 178
- disk pool, 63
 - adding partition, 69
 - availability, 67
 - cache, 68
 - create, 63
 - create from partition, 129
 - partition, 65
 - partition details, 69
 - read-only, 67
 - removing partition, 70
 - shared, 73
 - status indicators, 68
- disks, 123
- distributed lock manager (see lock manager)
- DNS cluster, 291
- domain, 35
 - create, 35
 - information, 36
 - properties, 37
 - storage, 37
- DRLM (see lock manager)
- E**
- e-mail
 - notification, 187, 316
 - SMTP server, 297
 - user, 225
- excluded tapes (pool), 96
- export pool, 114
- exporting configuration, 20

F

failback
 HA cluster, [288](#)
 shared disk pool, [86](#)

features
 new, [1](#)
 system, [3](#)

file access, [231](#)

file locks, [280](#)

file permissions
 active circle, [176](#)
 history, [176](#)
 Unix, [174](#)
 view, [176](#)
 Windows ACLs, [175](#)

file server, [259](#)
 access timeout, [262](#)
 activating, [260](#)
 activating shares, [268](#)
 protocol, [259](#)
 restarting, [261](#)
 statistics, [262](#)

filter
 administration log, [305](#)
 library inventory, [132](#)
 table views, [24](#)
 tapes in pools, [104](#)
 users and groups, [213](#)

fragmentation threshold, [98](#)

FTP protocol, [264](#)

G

group, [213](#)
 add user, [219](#)
 finding, [213](#)
 list, [213](#)
 managing, [218](#)
 share access, [220](#)

H

HA cluster, [285](#)
 add node, [287](#)
 create, [286](#)
 DNS, [291](#)
 edit parameters, [293](#)
 external software, [329](#)
 failback mode, [288](#)
 IP, [289](#)
 NetBIOS, [293](#)

 network settings, [295](#)
 NIC, [295](#)
 node election, [285](#)
 node priority, [288](#)
 OpenVPN, [292](#)
 restart, [293](#)
 settings, [288](#)
 shares, [295](#)
 types, [288](#)
 high availability cluster (see HA cluster)

I

import
 directory, [189](#)
 LTFS, [112](#)
 TAR file archive, [112](#)

inheriting values, [25](#)

integrity check, [109](#)

IP address, [51](#)

IP cluster, [289](#)

L

label, [127](#)

library, [131](#)
 configure devices, [327](#)
 detection, [131](#)
 information, [132](#)
 inventory, [132](#)
 localvars, [324](#)
 mailboxes, [109](#), [135](#)
 moving tapes, [134](#)
 nodes, [138](#)
 operations, [133](#)
 reinitializing, [134](#)
 SCSI address, [132](#), [138](#)
 synchronizing, [134](#)
 tape drive, [136](#)

license, [17](#)

local variables
 description, [323](#)
 encryption, [325](#)
 multipath, [326](#)
 network interfaces, [324](#)
 port ranges, [325](#)

lock manager, [277](#)
 advanced, [281](#)
 conflict resolution, [282](#)
 enabling on share, [243](#)
 list of locks, [280](#)

locks, [280](#)

log

administration, [303](#)

audit, [273](#)

levels, [53](#)

node, [52](#)

settings, [54](#)

statistics, [275](#)

LTFS, [93](#), [114](#)

import, [112](#)

LUN, [126](#), [127](#), [132](#), [139](#)

M

meta locks, [280](#)

MIB file, [300](#)

migrate files, [333](#)

retaining ACLs, [335](#)

robocopy, [333](#)

synchronizing, [334](#)

without ACLs, [334](#)

monitoring, [28](#)

N

NAC (see network access control)

NAS (see share)

NetBIOS cluster, [293](#)

network access control, [227](#)

hostname, [229](#)

IP address, [227](#)

shares, [230](#)

network interfaces, [51](#)

NFS connection options, [256](#)

NFS protocol, [265](#)

node, [39](#)

activate file server, [260](#)

activate shared tape pools, [49](#)

adding, [39](#)

adding to shared pool, [89](#)

cache, [45](#)

consolidation, [46](#)

failback mode, [86](#)

in shared disk pools, [84](#)

IP address, [51](#)

list, [27](#)

log, [51](#)

administrating, [53](#)

monitoring, [28](#)

moving, [40](#)

network interfaces, [51](#)

order of preference, [86](#)

parameters, [44](#)

performance indicators, [28](#)

properties, [51](#)

protocols, [259](#)

removing, [43](#)

restarting, [27](#), [42](#)

service management, [55](#)

sessions, [47](#)

shared access status, [85](#)

starting, [9](#)

stopping, [27](#), [42](#)

storage pools, [49](#)

version, [39](#), [44](#)

notifications, [311](#)

O

obsolete files, [239](#)

OpenVPN cluster, [292](#)

P

partition

adding local, [124](#)

adding shared, [126](#)

adding to a pool, [69](#), [79](#)

deleting, [129](#)

details, [69](#)

local, [124](#)

removing or deactivating, [70](#)

shared, [76](#), [124](#)

adding, [126](#)

group, [128](#)

inconsistent, [83](#)

removing or deactivating, [81](#)

reusing, [83](#)

status indicators, [79](#)

types, [123](#)

password

login, [14](#)

user, [217](#)

performance indicators, [28](#)

policy

archive parameters, [165](#)

archiving, [161](#)

association, [241](#)

constraints, [143](#), [164](#)

create, [142](#), [162](#)

data location, [145](#)

management, [283](#)

- processing, [236](#)
- processing nodes, [283](#)
- rules, [144](#)
- scheduling, [242](#)
- shares, [150](#)
- status indicators, [237](#)
- storage, [141](#)
- version management, [153](#)
- version management rules, [156](#)
- policy management, [283](#)
- pool
 - adding partition, [69](#)
 - availability, [67](#), [121](#)
 - cache, [68](#)
 - cloud, [117](#)
 - create, [117](#)
 - disk, [63](#)
 - create, [63](#)
 - partition, [65](#)
 - export, [114](#)
 - local, [63](#)
 - partition details, [69](#)
 - read-only, [67](#), [121](#), [121](#)
 - removing partition, [70](#)
 - shared, [73](#)
 - tape, [91](#)
- port ranges, [325](#)
- profiles, [171](#)
- protocol, [259](#)
 - authentication, [173](#)
 - CIFS, [263](#)
 - file access, [173](#)
 - FTP, [264](#)
 - NFS, [265](#)

Q

- quotas, [250](#)
 - rules, [251](#)
 - setting, [252](#)

R

- removing a node, [43](#)
- replication, [141](#)
 - inter-domain, [150](#)
 - schedule, [150](#)
 - enabling on share, [243](#)
- reset shared disk pool, [89](#)
- restarting a node, [42](#)
- robocopy, [333](#)

S

- SAN, [124](#), [131](#), [138](#)
- SCSI address, [132](#), [138](#)
- search
 - tapes, [104](#)
 - users and groups, [213](#)
- secondary tape pool, [93](#)
- service extension, [9](#)
- share, [231](#)
 - access controls, [244](#)
 - access rights, [221](#)
 - archives, [253](#)
 - details, [255](#)
 - files, [254](#)
 - filter, [256](#)
 - location, [254](#)
 - start manual archiving, [256](#)
 - connecting to, [256](#)
 - create, [232](#)
 - delete, [235](#)
 - depositing files, [258](#)
 - getting started, [231](#)
 - group, [233](#)
 - information, [237](#)
 - archiving policy status, [239](#)
 - policy manager, [238](#)
 - quotas, [239](#)
 - storage policy status, [239](#)
 - versioning volume, [239](#)
 - migrating files, [258](#)
 - move, [235](#)
 - node activation, [247](#)
 - parameters, [243](#)
 - policy association, [241](#)
 - policy indicators, [237](#)
 - policy processing, [236](#)
 - policy scheduling, [242](#)
 - quotas, [250](#)
 - renaming, [235](#)
 - settings, [240](#)
 - size limit, [250](#)
 - version management, [248](#)
- shared disk pool, [73](#)
 - access nodes, [74](#)
 - adding node, [89](#)
 - adding partition, [79](#)
 - create, [73](#)
 - election delay, [85](#)
 - inconsistent partition, [83](#)

- leader node, [84](#)
- managing nodes, [84](#)
- node access status, [85](#)
- node order of preference, [86](#)
- partition status, [79](#)
- partitions, [76](#)
- removing or deactivating partition, [81](#)
- reset, [89](#)
- reusing partition, [83](#)
- status indicators, [78](#)
- validate configuration, [89](#)
- SMTP
 - outgoing server, [297](#)
- SNMP Target, [299](#)
 - create, [300](#)
 - MIB file, [300](#)
- starting
 - active circle service, [9](#)
 - administration tool, [13](#)
- statistics, [269](#)
 - activation, [270](#)
 - file server, [262](#)
 - log file, [270](#)
 - contents, [275](#)
 - location, [271](#)
 - size limit, [271](#)
 - time, [272](#)
- stopping
 - active circle service, [9](#)
- stopping a node, [42](#)
- storage, [57](#)
 - devices, [60](#)
 - pools, [58](#)
- storage policy, [141](#)
 - constraints, [143](#)
 - create, [142](#)
 - data location, [145](#)
 - inter-domain processing, [150](#)
 - replication, [141](#)
 - replication schedule, [150](#)
 - rule, [144](#)
 - shares, [150](#)
- supervision, [311](#)
 - access, [311](#)
 - advanced, [320](#)
 - filtering, [313](#)
 - history, [314](#)
 - limit per share, [321](#)
 - list, [312](#)

- notes, [311](#)
- notifications, [316](#)
- registrations, [315](#)
- severity, [311](#)
- SNMP, [299](#)
- troubleshooting, [320](#)

T

- tape
 - adding, [108](#)
 - assigning, [107](#)
 - close, [109](#)
 - defragment, [111](#)
 - delete or remove, [108](#)
 - duplicate, [111](#)
 - fragmentation, [101](#)
 - import, [112](#)
 - information, [100](#)
 - integrity check, [109](#)
 - list, [58](#)
 - filter, [104](#)
 - locks, [280](#)
 - managing, [104](#)
 - moving, [134](#)
 - moving to/from mailbox, [135](#)
 - parameters, [100](#)
 - search, [104](#)
 - serial number, [100](#), [104](#)
 - unused, [111](#)
- tape drive, [136](#)
 - advanced information, [139](#)
 - cleaning, [137](#)
 - deleting, [137](#)
 - errors, [139](#)
 - maintenance, [136](#)
 - SCSI address, [132](#), [138](#)
- tape library (see library)
- tape pool, [91](#)
 - activating, [94](#)
 - adding tapes, [108](#)
 - assigning tapes, [107](#)
 - auto-expand, [98](#)
 - blank tapes, [95](#)
 - cleaning tapes, [96](#)
 - create, [92](#)
 - default, [95](#)
 - excluded tapes, [96](#)
 - export, [114](#)
 - fragmentation, [98](#)

- LTFS, [93](#), [114](#)
- mailboxes, [109](#)
- managing tapes, [104](#)
- pack, [93](#)
- parameters, [96](#)
- removing tapes, [108](#)
- secondary, [93](#)
- TAR archive, [93](#)
- types, [92](#)
- TAR archive
 - block size, [162](#)
 - import, [112](#)
 - integrity check, [109](#)
 - prefetch files, [243](#)
 - structured, [168](#)
- TAR archiving policy (see archiving policy)
- troubleshooting
 - advanced, [320](#)
 - disk pool status, [68](#), [78](#)
 - share policy status, [237](#), [239](#)
 - tapes in a pool, [102](#)

U

- user, [213](#)
 - access rights, [221](#)
 - add to group, [218](#)
 - authentication, [173](#), [174](#)
 - create, [214](#)
 - e-mail, [225](#)
 - finding, [213](#)
 - list, [213](#)
 - modify, [216](#)
 - password, [217](#)
 - profile, [217](#)
 - profiles, [171](#), [215](#)
 - administrator, [172](#)
 - sessions, [222](#)
 - share access, [220](#)

V

- validate shared disk pool, [89](#)
- version, [39](#), [44](#)
- versioning policy, [153](#)
 - adding to share, [248](#)
 - create, [154](#)
 - default, [154](#)
 - deleted files, [160](#)
 - rules, [156](#)
- views, [19](#)

X

- xcopy, [333](#)

